No One Drinks From the Firehose

How Organizations Filter and Prioritize Vulnerability Information

de Smale, S.; van Dijk, Rik; Bouwman, X.B.; van der Ham, Jeroen; van Eeten, M.J.G.

DOI
10.1109/SP46215.2023.10179447

Publication date
2023

Document Version
Final published version

Published in
Proceedings - 44th IEEE Symposium on Security and Privacy, SP 2023

No One Drinks From the Firehose: How Organizations Filter and Prioritize Vulnerability Information

Stephanie de Smale[1,2], Rik van Dijk[1], Xander Bouwman[2], Jeroen van der Ham[1,3], and Michel van Eeten[2]

[1]National Cyber Security Centre, The Netherlands
[2]Delft University of Technology
[3]University of Twente

2023 IEEE Symposium on Security and Privacy (SP) | 978-1-6654-9336-9/23/$31.00 ©2023 IEEE | DOI: 10.1109/SP46215.2023.10179447

***Abstract*—The number of published software vulnerabilities is increasing every year. How do organizations stay in control of their attack surface despite their limited staff resources? Prior work has analyzed the overall software vulnerability ecosystem as well as patching processes within organizations, but not how these two are connected.**

We investigate this missing link through semi-structured interviews with 22 organizations in critical infrastructure and government services. We analyze where in these organizations the responsibility is allocated to collect and triage information about software vulnerabilities, and find that none of our respondents is acquiring such information comprehensively, not even in a reduced and aggregated form like the National Vulnerability Database (NVD). This means that information on known vulnerabilities will be missed, even in critical infrastructure organizations. We observe that organizations apply implicit and explicit coping mechanisms to reduce their intake of vulnerability information, and identify three trade-offs in these strategies: independence, pro-activeness and formalization.

Although our respondents' behavior is in conflict with the widely accepted security advice to collect comprehensive vulnerability information about active systems, no respondents recall having experienced a security incident that was associated with missing information on a known software vulnerability. This suggests that, given scarce resources, reducing the intake of vulnerability information by up to 95% can be considered a rational strategy. Our findings raise questions about the allocation of responsibility and accountability for finding vulnerable systems, as well as suggest changing expectations around collecting vulnerability information.

1. Introduction

Contrary to popular belief, most attacks do not rely on advanced exploitation techniques or zero-days. Instead, they target known vulnerabilities which have been published months or sometimes years earlier [1]. Delayed mitigation for known vulnerabilities can be disastrous, as many high profile breaches have demonstrated. Think of Equifax [2], Maersk [3] and, more recently, the victims of the Microsoft Exchange attacks [4].

Against this backdrop, security experts urge organizations to mitigate the known vulnerabilities in their infrastructure. Public vulnerability disclosures—e.g., via vendor announcements, security websites, or vulnerability databases—are the necessary first step for the vulnerability management process [5]. Organizations might also discover vulnerabilities in their infrastructure in other ways, like pen testing and red teaming [6], but these are complementary and not a replacement for a continuous process of acquiring and evaluating information on newly published vulnerabilities.

The advice to mitigate known vulnerabilities sounds straightforward, until one takes the scale of the problem into account. The number of published vulnerabilities grows every year. For 2021, VulnDB reported over 29,000 vulnerabilities [7]. Furthermore, information remains dispersed. While platforms like VulnDB aggregate published vulnerabilities, none of the platforms are complete or always timely [8].

How do organizations cope with the flood of vulnerability information? In a year with at least 29,000 vulnerabilities, just assessing if a vulnerability is at all relevant for their enterprise infrastructure would already require over 500 evaluations every week—an enormous undertaking. This does not include the more complicated follow-up tasks of identifying all instances of the vulnerability in their assets, assessing the risks associated with these vulnerabilities, the impact of mitigation on business operations, and the prioritization, testing and deployment of patches or other mitigations.

Prior work has focused on the two different sides of this problem, rather than on how they are connected. On one side, there is a body of work that analyzes large aggregated vulnerability datasets [8, 9, 10, 11]. It studies patterns in discovery, timing and exploitation of vulnerabilities, often with the aim to support patch prioritization, since scarce resources and business impacts prevent mitigating all known vulnerabilities. On the other side, there are studies on the vulnerability management process of enterprises [6, 12, 13, 14, 15]. It looks at how system administrators acquire patch information, how they select which patches to deploy and how long it takes to deploy them.

There is an important gap between these two areas: how comprehensive is the vulnerability information that organizations seek out compared to the total set of published vul-

DOI 10.1109/SP46215.2023.00012

nerabilities? Li et al. [12] found that system administrators use various sources, such as vendor notifications, security advisories, blogs and forums. The authors concluded that we do not know how comprehensive or effective administrators are at acquiring the relevant information.

We aim to fill this research gap with an empirical study of 22 organizations in critical infrastructures and government services, where missing relevant vulnerability information is potentially catastrophic. Via the national CERT in the Netherlands, we contacted and interviewed a recipient of CERT security advisories in each organization. We transcribed and coded the interviews on how their organizations deal with the conflict between limited staff time and staying informed about their attack surface.

We first mapped where the organization allocated the task of acquiring vulnerability information. We then assessed what information sources the organization uses. The selected sources reflect large differences in the volume of vulnerabilities that the organizations are informed about, reflecting different points along a 'funnel' from the total input of more than 25,000 vulnerabilities per year to the output of a much smaller and manageable set of vulnerabilities that are actually evaluated by staff—and what 'manageable' means depends on the size of the organization. After analyzing these results, we interviewed the same respondents again about the consequences that their organizations suffered, if any, for their limited intake of vulnerability information. We make the following contributions:

- We present an intuitive descriptive model of how organizations reduce their intake of public vulnerability information via their selection of specific sources. The mix of sources reflects the volume that is manageable for them, even though this risks missing critical information about known vulnerabilities.
- We find that no organizations tried to acquire a comprehensive view of the flood of published vulnerabilities and none even reported using an aggregator like the National Vulnerability Database (NVD) or its commercial alternatives. Instead they implicitly or explicitly relied on curation by trusted authoritative sources to filter their vulnerability information. Prime examples are CERT security advisories, vendor notifications, security websites and vulnerability scanning solutions.
- We identify three trade-offs underlying the coping strategies of organizations: conducting independent assessments vs. trusting authoritative sources, acquiring information proactively vs. reactively, and relying on formalized vs. ad hoc organizational processes. There exists a relationship between the size of the organization and strategy chosen with these trade-offs, namely that larger organizations tend to be less dependent on trusting others, more proactive and more formalized.
- While the coping strategies dramatically reduce the intake of information, sometimes up to 95%, respondents expressed little concern for missing relevant vulnerabilities. This attitude is facilitated by the fact that no respondent could remember a security incident associated with missed vulnerability information.

2. Related Work

Public vulnerability disclosures are one step in an organization's vulnerability discovery process. There is an area of research related to the overall ecosystem of published vulnerabilities, e.g., around the NVD [16]. These studies analyze how to assess and prioritize vulnerabilities through the use of metrics and models. Notably, the CVSS metric [17] is often used to support decision making in filtering out the most critical software vulnerabilities, but according to its developers CVSS is actually unfit for this purpose [18]. Researchers have been critical of if such metrics accurately model risk. Allodi and Massacci [19] conducted an empirical study and found that using the commonly used CVSSv2 score to prioritize vulnerabilities was no better than doing it randomly. Spring et al. [20] stated that users of CVSSv3 are likely not learning from it what they are thinking they are learning from it: the score does not indicate risk but rather shows severity. The authors pose that the CVSSv3 formula has not been sufficiently justified, nor has it been empirically validated. They illustrate the CVSS metric is in practice "widely misused" for direct vulnerability prioritization instead of being one instrument among multiple inputs. This is congruent with a user study of CVSSv3 which found that scores are not consistently calculated, showing little difference in performance between the experts and laypersons [21]. Building on these critiques, authors have proposed to improve vulnerability prioritization by introducing new factors to the models [22, 23, 24], by developing alternative decision-making models [20] and by applying machine-learning to try to predict when a vulnerability will be exploited [25, 26]. However, metrics and models are just one side of the story.

There is a weak relationship between vulnerability information being made available and patching actually taking place. Research on organizational factors helps explain when and why remediation does take place. For example, from interviews with sysadmins, Li et al. [12] found compliance with organizational policies to be an important factor in when a patch is deployed. The authors discuss in depth the process by which sysadmins deploy updates, showing that they often used a staggered deployment model (i.e., separating development and production environments) which trades off timeliness for reliability. These authors further note that "[vulnerability] information is highly dispersed" and that "it is possible that some system administrators may lack the full coverage of relevant information", suggesting that this is "a nontrivial task for many". A similar example of organizational factors shaping outcomes is the study of Bouwman et al. [27] who found that when selecting sources of threat intelligence, organizations were not optimizing for network detection, as one might expect, but rather for their analysts' workflow by minimizing false positives. Alomar et al. [6] interviewed 53 security professionals and describe organizational factors in outsourcing detection within the vulnerability discovery process. They found that organizations that don't prioritize security often have a reactive stance towards detection of vulnerabilities. Building on this

study, we include public sources and third parties as crucial ways in how organizations learn about disclosed vulnerabilities and assess their relevance, and focus on tensions in organizational processes.

In short, research has highlighted the need of using a wide variety of sources in order to ensure coverage of vulnerabilities, but does not detail how this happens in practice nor at what scale. So prior work has not established how the two areas – the vulnerability ecosystem and vulnerability remediation by organizations – are connected to each other.

3. Methodology

We conducted two rounds of semi-structured interviews to study how organizations deal with the large volume of published vulnerabilities. Our study was approved by the Human Research Ethics Committee of Delft University of Technology. Respondents were explained in detail about the study, associated risks, use of information for which they provided informed consent. The conversations focused on public vulnerability discovery and triage.

3.1. Participant Selection

In collaboration with the national CERT, we approached 29 organizations in critical infrastructure sectors and central government that receive the CERT's official security advisories. As seen in Table 1, we contacted both large and small organizations. Organizations were emailed about the study via the email address that they have designated for receiving security advisories. We asked for one respondent who is responsible for receiving and processing security advisories and other vulnerability information. Two weeks after the first invitation to participate, we sent a reminder. If organizations did not reply after the reminder e-mail, they were contacted by phone.

In total, we conducted interviews in 24 organizations. However, 2 participants did not consent to having their interview data transcribed for analysis. We could therefore not include these interviews in our analysis, but the information we received in those interviews is consistent with the patterns we identified in the remaining interviews. So our analysis is based on 22 organizations, which means our invitation had a response rate of 76%. We interviewed respondents in 14 central government organizations and 8 organizations from critical infrastructure sectors, such as drinking water supply, energy, and finance.

3.2. Interview Protocol

We conducted four pilot interviews to validate the interview protocol and process. These interviews were not used in the data analysis. On the basis of these interviews, the protocol was refined. From May until June 2020, we conducted semi-structured interviews with professionals in the contacted organizations who were receiving and processing security advisories. The interviews were conducted in person or via video conferencing applications, lasting between 50 to 75 minutes each. Respondents were volunteers and were not compensated for the interview.

Since we recruited participants who receive National CERT security advisories, we used these advisories as a starting point for eliciting information about the organizational work flow around seeking and receiving vulnerability information. Analyzing the role of public information in the vulnerability discovery and triage process, the interview protocol focused on two topics: the inflow of public vulnerability information (acquisition), and processing, prioritizing, and reviewing vulnerability information from different sources (triage) [5]. Participants were asked to complete a short post-interview questionnaire, about their role, years of experience as a practitioner within the organization, as well as demographic questions about the size of their organization, size of their incident response team, and sector.

After a preliminary analysis of the results from these interviews, we found that no organization was ingesting anywhere close to comprehensive vulnerability information. This finding was troubling and raised follow-up questions. Did the organizations consciously design their limited intake of information? Did they ever suffer consequences for missing vulnerability information? To answer these questions, we conducted brief follow-up interviews in February 2022. Of the 22 original respondents, 16 agreed to participate in the second round. The respondents received a summary of our preliminary findings and were then asked additional questions on source selection, incidents based on missed information, and evaluation of their selected sources. All interview questions are included in Appendices B and C.

3.3. Interview Coding

Data collection and transcription was done by the two lead researchers in collaboration with two other researchers. Data analysis was done the two lead researchers using inductive thematic analysis [28], which has distinct phases: *i) Familiarization with the data*. Actively reading and re-reading the data and summary, becoming immersed and intimately familiar with its content. *ii) Initial Coding*. The two lead researchers coded the interviews in Atlas.ti [29], independently verifying the results. Results generated 413 initial codes. They summarized, grouped, and downsized the preliminary results into 6 code clusters of 220 codes. *iii) Theme generation*. Lead researchers discussed broader patterns based on results and initial coding groups. Three potential themes were defined in a group discussion with two other researchers, which were explored using network analysis in Atlas.ti. The 147 codes were mapped in three thematic networks, drawing semantic linkages between them. *iv) Reviewing themes*. The lead researchers independently verified results in two sessions. This is a suitable way to ensure reliability of qualitative findings [30]. Themes were refined into trade-offs using 79 codes (see table 4) and overlapping issues were mapped. *v) Defining themes*. Lead researchers described the sources used by organizations and the strategies used to filter information, writing out the

narrative in the result section. They selected representative quotes and added them to the write-up. *vi) Report findings*. Finally, results were contextualized in relation to existing literature.

4. Organizational Roles

We first set out to describe where vulnerability information enters the organization. Where do organizations allocate the responsibility to receive and process this information? We find that there is a large variety among organizations. Various professionals, ranging from analysts, system administrators, to management, are assigned the task to deal with vulnerability information. Organizations need to handle this according to the means and staff available.

While some of our respondents work in organizations with thousands of employees and large incident response teams, there are organizations in critical sectors with less than 500 employees, relatively small IT departments and even smaller incident response teams. The overview in table 1 illustrates that there is great variance in the make-up of these organizations.

Table 1 provides an overview of the organizations. They are very different in terms of total size (FTE). We interviewed 6 respondents from large organizations with more than 5000 employees, 8 respondents in medium-sized organizations with 501-5000 employees, as well as 8 respondents employed in smaller organizations with 1-500 employees. Respondents worked in different areas, some are specialized, such as Security Operations Centers (SOC). Other respondents worked in security or IT-security departments or in general IT departments. Finally, several interviewees worked in business operations.

Table 2 provides an overview of security advisories recipient roles. Striking is the differentiation in managerial levels, from the CISO level to CISO office staff to security specialists to general system administrator roles. Some recipients of security advisories worked as an analyst responsible for monitoring vulnerability information for a whole ministry. Others were system administrators working for a small government organization with a part time role as a security officer.

The majority of respondents with analyst roles are employed in Security Operations Centers (SOCs). Respondents working in a SOC worked at medium or large organizations. We observed a more formalized division roles in these SOCs. Dedicated security analysts receive and analyze vulnerability information before sending it to stakeholders within the organization. All SOCs use a functional mailbox that is accessible by multiple team members and other stakeholders within the organization to receive and process vulnerability information. All SOCs but one structurally use vulnerability scanners as a source of information in conjunction with other sources, described in detail below. To process these extra sources of information, we see that organizations with a SOC have at least medium size IR teams, regardless of the total size of their organization. These respondents are part of organizations with larger resources to process vulnerability information.

Interviewees in three organizations had dual roles within the organization; specifically all dual roles were employees who combined the work of security officer with that of system administrator. Whereas the total organization size is diverse, the IR teams were all small. As opposed to the formalized division of roles at SOCs, the dual roles have a much less formalized position. These roles both receive information and are responsible for assessing and the remediation of possible vulnerabilities, fulfilling tasks other organizations employ several people for in various roles.

Important to note is the unequal distribution of capacity to monitor vulnerability information. On the one hand there are large teams with specialized analyst roles who are able to monitor different sources full time. On the other hand, there are individuals who have this role part time, aside from other, non-security tasks. Furthermore, it provides information on where responsibilities for processing and evaluating vulnerability disclosures lie. Where one organization has a team to assess information, others are solely responsible for this task. But the stakes for each of these organizations are the same; they are all important with respect to critical infrastructure services and national security.

5. Sources of Information

How do organizations learn about public vulnerability disclosures? We asked the respondents to give an overview of the sources. We coded the answers (Table 3). In one sense, the most important answer is the one missing from

TABLE 1. PARTICIPATING ORGANIZATIONS

#	Size	Department acquiring vulnerability info.	IR team size	Conducts structural scanning	Filtered advisories
1	Large	SOC	5-10	No	Yes
2	Large	CISO office	11-25	Yes	No
3	Large	Technical support	0-5	No	No
4	Large	SOC	25+	Yes	No
5	Large	SOC	25+	Yes	Yes
6	Large	Info. management	0-5	No	No
7	Medium	Business operations	11-25	Yes	Yes
8	Medium	SOC	11-25	Yes	Yes
9	Medium	CERT	5-10	No	Yes
10	Medium	Technical support	5-10	Yes	Yes
11	Medium	Security	5-10	Yes	No
12	Medium	SOC	5-10	Yes	No
13	Medium	Technical support	25+	Yes	No
15	Medium	IT Security	5-10	No	Yes
15	Small	Technical support	0-5	No	No
16	Small	IT	0-5	No	Unknown
17	Small	Security	0-5	No	Yes
18	Small	Info. management	0-5	No	No
19	Small	Business operations	0-5	No	Yes
20	Small	Security	0-5	No	Yes
21	Small	Technical support	0-5	No	Yes
22	Small	Operations	0-5	No	Yes

Organization size: small=1-500 FTE, medium=501-5000 FTE, large=5000+ FTE. (Ranges are used to protect anonymity)

TABLE 2. RESPONDENT ROLES AS RECIPIENTS OF SECURITY ADVISORIES

Roles	Respondents *n=22*
CISO	4
Dual role: SysAdmin/Sec Officer	3
Security Analyst	3
System Administrator	2
Advisor IT Network	1
CISO Office	1
Network Adminstrator	1
Operational Security Lead	1
Privacy & Security Specialist	1
Risk manager	1
Security Engineer	1
Security Officer	1
SOC Analyst	1
SOC Team Lead	1

TABLE 3. SOURCES MENTIONED BY RESPONDENTS

Sources	Respondents *n=22*
Security advisories	22
Vendor notifications	19
Security websites	16
Vulnerability scanners	9
Threat Intel Platforms	7
Social media	5
International CERTs	3
Professional communities	3
OSINT monitoring	2
Personal contacts	2
External SOC	2

the table: no organization is monitoring the stream of over 25,000+ vulnerabilities that are published by databases such as VulnDB, let alone the information that is also missing from these aggregators [8]. The sources that the respondents mentioned all imply a lot of information has already been left out before it is even considered by the organization. On average, respondents mentioned they use 4 sources, but the variance is significant: 4 respondents mentioning using only 2 sources, while one respondent mentions 7. The list is largely in line with the findings of Li et al. [12], who interviewed only system administrators. We will discuss the four most prevalent sources and relate these to the characteristics of the organizations.

Security advisories are received by all respondents, which reflects that their organizations were selected on the basis of them receiving National CERT security advisories. The advisories represent a heavily filtered set of vulnerabilities. The National CERT sends around 1200 advisories each year for vulnerabilities that it has selected after a risk assessment of likelihood and impact. This volume constitutes less than 5% of the total number of published vulnerabilities in 2021. Interviewed organizations can opt to receive only specific CERT advisories based on their IT assets. As shown in Table 1, 12 of these 22 organizations have chosen to receive only a subset of these advisories. We did not find clear indicators for why organizations choose to receive a subset or not.

Relying on the advisories of the National CERT reduces the amount of security advisories recipients must process every week, but it also indicates a higher degree of trust in a third party to pre-select advisories that are important to the recipient. For organizations mainly relying on the CERT advisories, this implies a filtering out of over 95% of all vulnerability information that is never received, let alone evaluated for relevance.

Vendor notifications are the second most frequently mentioned source. They are typically made up of periodic lists with vulnerabilities sent by software and hardware suppliers. For vulnerabilities that are deemed more critical by the vendor, organizations receive a notification and a possible path of action. One respondent described the outreach of vendors (13): *We have a fair amount of vendors who share information whenever there is a vulnerability. So I'd also consider that a source.* The time between lists differs per vendor. Some send out weekly notifications, while others send out a list each month.

In some cases, respondents directly communicate with the vendor on a specific notification. Organizations value the additional information vendors can provide on a specific product and the speed of the notifications, but they are rarely seen as a full replacement for security advisories. Moreover, many vendors or open source projects have no direct channel to the users or admins of their products, so they rely on public channels to reach their user base. This means direct vendor notifications will only provide organizations with a fraction of all vulnerability disclosures.

Open source security websites are frequently used by many respondents, varying from national, international security news, as well as expert blogs or fora. Said respondents 20: *Yeah so these are mainly technical news sites: [...] Hackernews, Securityfocus, Threatpost, Collegeblog, Opensourcesecurity.*. Social Media were mentioned to a lesser extent: *Within our team we're active on Reddit to keep up with things... Different security Twitter accounts and so on. Often it's issues that are posted within the community. Then we know when something starts popping up around a specific product that we have to keep our eyes and ears open*(13).

Eight respondents relied solely on monitoring websites in combination with security advisories and, sometimes, vendor notifications. These organizations all have a total size of fewer than than 500 employees and incident response teams of less than 10 people. This could indicate that the choice to rely on these two or three types of sources is driven on the limited capacity and resources of the organization. The amount of vulnerabilities covered varies per website, but these tend to be vulnerabilities which are more newsworthy and are unlikely to contain more than a few vulnerabilities per day. This means that organizations learn about less than 5% of all published vulnerabilities via this channel – and a significant portion of this set is likely to overlap with the security advisories. Overall, these practices imply a large amount of trust in the third parties behind these sources, assuming that they will report on all relevant vulnerabilities for an organization.

Vulnerability scanning services like those offered by Qualys and Rapid7 Nexpose automate the process of identifying what vulnerabilities are present in organizations. Nine organizations use vulnerability scanners in-house and two have outsourced scanning to external SOCs that monitor vulnerabilities for them. All respondents that use vulnerability scanners employ at least 500 people. Also in this category, we see the use of a number of other automated types of sources such as external SOCs that supply the organization with vulnerability notifications. In a large enterprise environment, these services can scan for thousands of different vulnerabilities and their often voluminous results basically represent the intake of a large amount of vulnerability information – in fact, the largest amount of all sources mentioned by the organizations. That said, there is still a significant degree of trust implied by relying on these service providers, since the organizations assume that the services capture all relevant vulnerabilities. As respondent 11 put it: *Nexpose automates this [vulnerability scanning]. At six o'clock it receives an update and if there is a CVE-score of 8 or higher then it scans for the vulnerability directly, once per day. Some zones are scanned broadly, others are scanned once per day*. It is not transparent how comprehensive these scanners are but it is virtually impossible for them to be complete, as not all vulnerabilities can be scanned for automatically. Plus, some vulnerabilities are associated with protocols that are not supported by the scanners or the networks that they are deployed in. Organizations do not appear to have any sense of how comprehensive these services are, nor do they seem to worry relevant vulnerabilities—which reflects an implicit trust in the providers of these services.

Threat intelligence platforms (TIP) provide an interface to commercial feeds of information about adversary behavior, as well as open source blocklists [27]. These sources are not focused on vulnerability information directly, but reporting on adversary behavior informs the recipient which vulnerabilities are currently being targeted by attackers. Four organizations indicated that they made use of threat intelligence as a service, receiving e-mail notifications or periodic overviews. Two of these respondents have access to automated tooling. These are also part of teams with at least 25 people. The other two had a much smaller capacity and use vendor notifications and open source websites, in addition to the information managed by Managed Security Service Providers (MSSPs).

Even for a medium-sized enterprise, TIP services can lead to hundreds of notifications in their mailbox. Said respondents 19: *On average each day I receive maybe 20 or 30 notifications by mail? Maybe I'm forgetting some, because not all of them are important*. How organizations choose to manage threat intel information matters. On the one hand, opting for paid services to curate and filter relevant information quickly diminishes the amount of information that needs to be assessed. On the other hand, monitoring events can quickly increase the amount of notifications to assess. In the end, only a very small fraction of vulnerabilities are observed to be attacked in the wild [31, 32], and only these vulnerabilities would show up in threat intelligence.

6. Coping Strategies

The different sources reflect varying degrees of comprehensiveness in terms of vulnerability information. As noted above, no organization is trying to comprehensively track this information, not even in the reduced format of aggregation platforms like VulnDB or NVD. Every organization has to some degree reduced their intake of vulnerability information, even though this implies accepting the risk that they miss learning about vulnerabilities that could be present in their infrastructure. This is basically a coping strategy that seeks to balance the need to learn about known vulnerabilities without overwhelming the organization by more information then it can process and evaluate.

From our coding of the interviews, we identified three interrelated trade-offs that shape the coping strategies of organizations and, thus, the degree of comprehensiveness of the vulnerability information that they acquire and further triage towards actual mitigation and patching actions. The overview provided in Table 4 illustrates key elements linked to each trade-off, followed by the main codes with which we labeled the interviews. The trade-offs are sliding scales rather than dichotomies. The first trade-off is between trusting authorities and third parties to select relevant vulnerability information versus independently seeking and evaluating vulnerability information. The second trade-off is between proactive monitoring of public vulnerability disclosures to enable curation and assessment of information versus reactive responses upon receiving vulnerability information sent by outside parties. Third, we see a trade-offs between adopting highly formalized processes with automated tools to acquire, structure and prioritize information versus ad hoc processes that rely more on people and manual labor to acquire and evaluate information in their triage process.

6.1. Independence Versus Trust

The first strategy considers to what extent organizations rely on others to learn about new vulnerabilities – in other words, trust placed in third parties.

6.1.1. Independent from Others. Independence increases if organizations source and assess a multitude of sources and cross-references them to formulate risks. For instance, by not relying on one source, but by collecting sector-specific advisories, national and international advisories and comparing them. As shown in Table 4, independence means relying less on national authorities or third parties for risk assessments, and being able to process information and make assessments about vulnerability information in-house. One organization used various information sources to create its own security advisories, tailored to the needs of its IT-infrastructure. As a respondent explained: *You have to make your own judgement [...] That is why you need checks and balances, what is said internationally about this vulnerability? Then we look at communities, friends and acquaintances we trust very highly. How did you assess that [vulnerability]? (2)*

6.1.2. Trusting other parties. Organizations that have a high degree of trust in other parties, rely more on these parties to provide and assess vulnerability information. This is seen in how some organizations are likely to adopt assessments by authoritative peers and third parties, or outsource (parts) of its vulnerability management capability altogether. Organizations with a high degree of trust outsource selecting—i.e., filtering out—vulnerability information. This happens through authoritative peers such as national CERTS, but also vendors play a key role. As seen in 11 organizations, outsourcing the supply of vulnerability information is in some cases necessary. The organization does not have the capacity to organize and process the complete information stream, making them dependent on external capacity for monitoring vulnerabilities.

The national CERT adds a risk assessment to their security advisories. They evaluate whether the risk is "low", "medium", or "high" along two dimensions: the likelihood of a vulnerability being exploited; and the potential impact when the vulnerability is exploited. Especially the risk scores given by trusted third parties are appreciated by these organizations, as it allows security teams to better advocate for mitigating measures within the organization. One respondent stated that the risk score from the national CERT determined further handling: *When the National CERT says that a vulnerability medium/high [impact/likelihood] and we base further steps on that. So that's leading. Sometimes we have questions, like: How did you come to that medium/high? Then you have interesting conversations on the topic – but the national CERT is leading* (17). Security advisories sent by national CERTS play a role in creating urgency: *Alarm bells go off when its a High/High* (18). In Table 4, key codes for trusting organizations center around reliance on external stakeholders for triage.

In sum, the coping strategies of organizations are distributed along this scale going from independence to trust. Organizations that are highly trusting on third parties tend to rely more heavily on CERT advisories as their main source of information, even though these filter out at least 95% of the published vulnerabilities. Particularly smaller organizations tend to be higher on trust. This pre-selection done by the National CERT is meant as an aid, but not replacement of independent evaluation of information.

At the other side we see mostly large organizations that make more independent evaluations based on scanning service outcomes (section 5). Their vulnerability information is more comprehensive with more automatic triage. This does not mean that these organizations are completely independent, as they still rely on the curation by the scanning services.

6.2. Proactive Versus Reactive

The second trade-off revolves around whether organizations have ongoing processes to seek and evaluate vulnerability information or if their triage process starts after they are informed of a vulnerability by a external party.

6.2.1. Proactive Activities. Proactive strategies for prioritizing information seek to detect vulnerabilities and identify risks, to collect and share information as a continuous process, with related codes shown in Table 4. A proactive approach means organizations determine whether publicly disclosed vulnerabilities are applicable to their organization [5]. Respondents periodically scan network segments to determine risks. Medium, severe and critical vulnerabilities are given a deadline to resolve, whilst low ones are registered but are remediated in normal patch cycles. Said a respondent working in a medium-sized organization cross-references sources: *A vulnerability scanner just monitors the CVE list. But its satisfying when you can look it up manually and check in your scanner if its correct* (11).

Nine of 22 respondents use in-house detection to analyze IT-infrastructure for vulnerabilities. All interviewed organizations that employed scanning software regarded them as an important source of vulnerability information.

Another proactive activity is to combine vulnerability information with threat information such as indicators of compromise (IoCs) or Tactics, Techniques, and Procedures (TTPs). Technical information of possible exploits are structurally gathered by (specialized) team members. Essential information is extracted from third party notifications and incorporated in these self-created bundles, combined with information of their respective environments and/or products. This is then sent to stakeholders within the organization. The ability to process and assess CVE information gives organizations a time advantage in assessing the risk of a certain vulnerability.

6.2.2. Reactive Activities. Reactive approaches take as their starting point assessments made outside the own organization. Such a process is not a continuous process, but dependent on outside triggers with codes shown in Table 4. As an example, advisories on critical vulnerabilities set in motion processes to assess and then resolve the vulnerability. An interviewee stated: *We receive a call [from the National CERT] when it's a High/High advisory [impact/likelihood]. So we have a big incentive if it's really critical [...] Once I see a High/High, I drop all other work and I'll first determine the impact specifically for our organization* (20).

Of the 22 organizations, 6 had security teams with five or fewer members. These organizations mentioned reactive strategies. None of them employ scanners or another form of automated vulnerability tooling. All sources of vulnerability information, most importantly advisories from suppliers and CERTs, are received by mail. The security teams focus on processing these advisories and possibly check information through open source news sites before advising system administrator teams and management on possible actions.

Instead of processing IoCs, reactive organizations rely on third parties such as vendors and government agencies like CERTs to deliver vulnerability information to their teams. This information is packaged with analysis, patch information and/or mitigating measures and risk scores to indicate follow-up actions for organizations. Respondents

TABLE 4. SUMMARY OF TRADE-OFFS

Trade-off		Observed behaviors	Frequency
Assessment	*Independent.* Organization assesses vulnerability information in-house.	Do not rely on national authorities or third parties to prioritize information flow.	4
		Independent of assessment third parties on vulnerabilities.	12
		Able to weigh and counter-balance prioritized information by external parties.	8
	Trusting. Organization relies on authoritative peers or third parties to assess and filter information.	Trust in third parties to process less refined sources of information and deliver relevant results to act upon.	12
		Trust in third parties to be notified proactively on current threat levels of vulnerabilities and steps to take to remediate them.	14
		Risk assessment by national authority driver in urgency of remediation.	10
Collection	*Proactive.* Actively search for vulnerability information.	Capacity to analyze less contextualized sources of information (e.g. IoC's, CVE's).	9
		Capacity to filter relevant vulnerability information for organizations IT assets.	8
		Conduct periodic scans of IT-environments to determine risks.	11
	Reactive. Information triggers activity.	Limited capacity to filter relevant vulnerability information.	9
		Follow up actions reliant on determined risk levels by third party.	14
		Outsourcing parts of triage process.	8
Processes	*Formalized.* Explicit organizational structures.	Strong division of roles and tasks.	12
		Automated triage process.	4
		DevOps/Agile work environment.	4
	Ad hoc. Informal workflows.	Weak division of roles and tasks.	12
		Manual curation of information.	18
		Gathering of information based on personal interests.	8

rely on third parties for these parts of the triage process as well, not just for the acquisition of information.

The trade-off between proactive versus reactive activities is closely related to vulnerability discovery activities within the organization. Alomar et al. [6] described activities such as penetration testing, blue teaming and red teaming. Where these authors found reactive attitudes in organizations with security lower on the priority list, we found these reactive activities to be associated with a limited capacity for vulnerability triage.

6.3. Formalized Versus ad hoc Processes

A third trade-off we observed is around formalization. Some organizations prefer to formalize their processes and develop tools to structure information flows about vulnerabilities. Other organizations prefer organic, ad hoc processes where professionals manually curate information flows and often use personal communication to manage vulnerabilities. This trade-off is closely related to proactive and reactive approaches, but not exactly the same. An informal process managed by a professional can still be proactive, but this is more difficult, since it doesn't scale. Large information flows are more likely to overwhelm the capacity of a small number of professionals working organically.

6.3.1. Formalized Processes. Organizations that formalize processes have a clear division of roles and tasks. They use and develop tools to structure information flows. Table 4 shows codes from the interviews such as formalized process and separation of roles to describe this strategy. An interviewee stated that: *We differentiate between handler groups. I have the coordinating role for handler group privacy and security. So I monitor if things are prioritized enough and if it runs smoothly* (22).

Another code associated with formalization is "high capacity because of automation". Formal structures are supported through the use of tools. One clear example which is widespread is the use of ticketing systems where notifications are loaded into, preferably in automated way. These systems are intimately connected to the formalized process of handling vulnerability information. Critical vulnerabilities are given the high priority label. This tooling might be actively developed. Several organizations mentioned having DevOps capabilities in their IR team. As one respondent from a medium-sized organizations shared: *In principle, we have three teams within our small team. We have the engineering team, which me and another colleague are a part of. We're mainly preoccupied with automation, so that we decrease repetition of labor. We have the security monitoring team to bring that to a higher level. And we have the tooling team which is preoccupied with technical and functional maintenance of our tools, and life-cycle management if needed.* (8)

One more indicator of formalization is that vulnerability advisories are received in functional mailboxes instead of personal mailboxes. This then requires a process to make sure that the mailbox is regularly reviewed and information is not missed. As one respondent put it (8): *The mailbox is monitored daily. Everyday we assign a man of the day*. Sharing information inside and outside the organization is also automated. Two respondents had structured and automated processes to gather relevant information for their infrastructure and spread it to stakeholders within their organization, such as the respective system administrators. A respondent working in the SOC of a large organization describes how they define use cases with specific key words and aggregate them from a variety of public sources.

6.3.2. Ad hoc Processes. Organizations which have ad hoc processes in their vulnerability triage use personal communication and manually curate information flows. One element of ad hoc processes, as shown in Table 4, is the weak division of roles and tasks. Individuals often have multiple roles within the triage process, particularly in teams with fewer than five FTE. Said respondent 10: *I'm an Information Security Officer, as we call it here. Next to my role as system administrator. Because of this, I am the primary person preoccupied with the vulnerabilities I see in [Windows Defender] ATP*. They are security practitioners responsible for monitoring vulnerability information, as well as system administrators responsible for patching the vulnerability. The personal communication works because the organizational hierarchy is relatively flat and individuals can be contacted easily.

Table 4 shows another element of ad hoc processes, namely manual curation of information. Receiving vulnerability information in personal mailboxes keeps the information ready at hand for the practitioner, which is particularly beneficial for individuals with dual roles. This works if the receiver of information is not overloaded and is knowledgeable on the assets of the organization. In asking how a respondent assesses if the information is relevant, they answered: *From the knowledge [on assets] I have. If I don't have it, I'll Google it. If I can find anything about the software we use I tackle it to look on our network. And if I think we don't have that software I'll write it down as well* (15). These informal processes are also more sensitive to what specific people deem most urgent. As one interviewee described this: *On the basis of appetite, interest and time of the analyst, because when we are busy with design or development, you see that the focus shifts towards that, so yeah then we're less focused on that* (12).

Both side of the trade-off have their pros and cons. The formalized approach is rigid, comes with high overhead and can mean that responsibility for correct triage of vulnerabilities is dispersed across various specialized roles. On the other hand, the informal ad-hoc approach does not scale very well, so a lot of vulnerability information might be missed, and the effectiveness of triage is dependent on specific individuals.

7. Learning About Known Vulnerabilities

The interviewed organizations in critical infrastructures and government exhibit coping strategies to reduce their intake of vulnerability information. Each of them has adopted a certain mix of sources that reflect underlying trade-offs around trust, proactivity, and formalization. Different sources represent different orders of magnitude in terms of the number of vulnerabilities the organization learns about.

We can visualize these sources, and the coping strategies that they are a part of, as points along a funnel (Figure 1). The input of the funnel is the total pool of newly published vulnerabilities. In 2021, this contained at a very minimum over 29,000 vulnerabilities—given that VulnDB, like all aggregators, is not complete [8]. The output of the funnel

TABLE 5. ORGANIZATION SIZE VS. VOLUME OF LARGEST SOURCE OF VULNERABILITY INFORMATION (ORG. NO. FROM TABLE 1)

Org. size (FTE)	Volume of largest source (vulns/year)		
	10,000	1,000	100
Large (>5000)	2, 3, 4, 5, 6	1	-
Medium (>500)	7, 8, 10, 11, 12	9	-
Small (<500)	-	15, 16, 17, 18, 19, 20, 21, 22	-

is the number of vulnerabilities that enter the organization's triage process. In between are various sources, ordered by the order of magnitude that they roughly correspond to in terms of how many vulnerabilities the organization learns about from that source in the course of one year. Of course, in practice there is a lot of variance across sources, even of the same type, so our placement along this dimension reflects a rank order rather than a numerical scale.

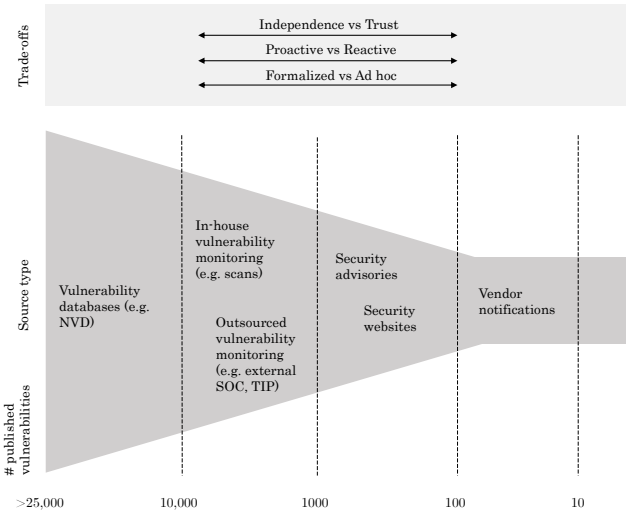

Figure 1. The 'funnel' of vulnerability information, a descriptive model of coping strategies where organizations restrict the yearly volume of vulnerability information they acquire through the selection of sources.

On the left side, we find the closest approximation of the total set of published vulnerabilities: the aggregators like VulnDB, NVD, etc. They capture information on tens of thousands of vulnerabilities. Next, we find the vulnerability scanning services. Rather than acquiring and processing comprehensive vulnerability information themselves, organizations rely on the providers of these services to do so. Organizations take the results from the scans as the vulnerabilities they need to triage. It is not clear to how comprehensive these services are, but respondents reported learning about thousands of vulnerabilities from running these tools across their infrastructure.

One further step to the right, we find sources like threat intelligence and vulnerability monitoring by MSSPs and external SOCs. These focus more on vulnerabilities that might be targeted by adversaries and that the client thus needs to learn about. They are less comprehensive than

dedicated vulnerability scanning services that aim to find vulnerabilities across the whole infrastructure. The next step down the 'funnel' brings us to sources like security websites and blogs. This was reported to result in a few vulnerabilities per day at most. Finally, the smallest source that most organizations used were the vendor notifications. This quantity is of course highly dependent on the type of vendor and the number of vendors, but few vendors report new vulnerabilities more than a few times per month.

All organizations reported using multiple sources alongside security advisories, between 2 and 7. Their overall intake of vulnerability information is driven by the largest source that they use. So for the 9 organizations that reported using scanning services, the funnel is rather short and much less narrow than for 11 organizations whose largest source is the CERT advisories. The longer the funnel, the more information about known vulnerabilities has been filtered out along the way. An organization that relies primarily on security advisories and vendor notifications will only learn about just a few percent of newly published vulnerabilities.

We approximated the intake of organizations by asking how many vulnerability messages the organization receive. Some respondents mentioned between 5 and 10 messages a day from various sources. They complement the advisory source with vulnerability information from vendors once or twice a month. Other organizations reported receiving more than a 100 messages each day.

The position along the funnel reflects the three trade-offs that we discussed in the previous section. The more to the right, the more trust the organization – implicitly or explicitly – places in the entities that provide this source. The unspoken assumption is that the entity is responsible for identifying the relevant new vulnerabilities for the organization. Some respondents explicitly stated that this is what they expect from the national CERT, even though the CERT does not intend its advisories to be used in that way. It does not claim to identify all relevant new vulnerabilities and sees the recipient as ultimately responsible for acquiring this information. A position to the left of the funnel is also associated with more proactive collection of information and with more formalized processes, while to the right we find more reactive strategies and ad hoc processes.

We find a clear relationship between the size of the organization and the size of the sources. Table 5 shows organizations categorized by size (large, medium and small) and by the largest source of vulnerability information that they reported using. The separation between large to medium organizations on the one hand and small organizations on the other hand, is clear. Larger organizations ingest an order of magnitude more vulnerability information than small organizations. They use in-house or outsourced monitoring capabilities, while small organizations rely more on security advisories and open source security websites as information sources, processing lower amounts of vulnerabilities. This relationship between size and sources suggests that organizations reduce the intake of information to not overwhelm their capacity, rather than developing the capacity needed to process information about new vulnerabilities.

8. Revisiting Intake Strategies

From our first round of interviews, a set of practices emerged that stand in contrast to the assumption of many experts that organizations need to comprehensively acquire vulnerability information in order to remedy known vulnerabilities [12]. Where best practice according to CISA is that organizations should match sources of vulnerability information with their assets, so each unique asset has a dedicated list [33], NIST emphasizes the importance of comprehensive sources to assess the risk of vulnerabilities to the organization [34] and industry vulnerability management models focus on automating and standardizing intake of information [35, 36]. But crudely put: none of our respondents' organizations acquired even close to comprehensive information. Organizations with fewer people and tools to process sources of information, choose to more severely limit their information intake. One such coping strategy is to rely fully on vulnerabilities categorized as critical by trusted external parties in security advisories. Respondent 3 explained: *If they [the national CERT] call me outside of office hours about a Windows patch, I know that they have rated the vulnerability as 'critical' and that I should better respond adequately and make time in my schedule to apply it soon*.

These best practices raises two follow-up questions. First, are the organizations conscious about how much information they are missing? Have they deliberately chosen and evaluated their set of sources or has this set organically grown over time? Second, did they ever experience an incident associated with a known vulnerability that they were not informed about? In other words, did they suffer negative consequences for reducing, sometimes drastically, their intake of information?

To collect answers to these questions, we contacted the same 22 respondents again for a short second interview. We were able to speak with 16 of them (73%). We asked them *i)* whether their set of sources was the outcome of a formal process or whether it grew organically over time, *ii)* if they had the idea they were missing relevant vulnerability information, and *iii)* if they had ever suffered an incident around a vulnerability they were not informed about. If yes, we asked if this caused changes in their set of information sources. If no, then we asked if there had been any changes to these sources in the past two years and, if so, for what reasons (the protocol is included in Appendix C).

Four respondents (25%) reported having a formal procedure for periodically evaluating their vulnerability notification sources in terms of relevancy. Most organizations said the sources had grown organically. For example, when a product is purchased from a new vendor, this often includes receiving vulnerability information on that product as a service. This source is then incorporated in the process. In the words of one interviewee: *To be completely fair, we consist of a limited number of network and system administrators. When you mention [sources], you'll notice it is not a conscious decision.* (12)

When asked if sources were added as a deliberate choice, respondents mention two types of sources that were added as a deliberate decision to their sources: threat intelligence services and vulnerability scanning services. Because these tools require a paid subscription, capacity and resources to effectively incorporate in the existing processes, the choice to add them to their sources, organizations use a review process to see which service best fits their needs. Vulnerability scanners require a degree of formalization of the vulnerability information intake process, because of the volume of information they produce, which requires the organization to free up capacity. The deliberate choices for these services is opposed to vulnerability information a vendor provides when purchasing a product or service, since this information is generally sent in the form of an e-mail, the information can easily be incorporated in the existing process without changes. As respondent 16 put it: *Sources come and sources go. The only processes that we formally arranged are the National CERT advisories and the scans we conduct.*

Surprisingly, ten of our respondents were confident that through their current sources, they were not missing relevant vulnerability information. They expressed being satisfied with their selection of sources. As one respondent stated: *I don't think we have [missed vulnerability information]. Until now we never had an incident that made us question our scope* (13). Of the respondents who felt less confident, two stated that they were positive that they had missed vulnerability information. The other four respondents stated various reasons why they felt they missed vulnerability information. Two respondents worried less about the offering but more about the timeliness of information they receive. They would like to receive vulnerability information before it is made public. One respondent said: *"I don't think anyone can claim to not miss anything. As I said earlier, we have a lot of products and some we don't even know we have. The Log4J case showed clearly, you can have an overview of what you have but tomorrow you discover you miss a lot."* (16)

The satisfaction with the existing, limited intake of information seems directly related to their answer to the next question: did they ever suffer an incident that was associated with missing vulnerability information? Of 16 respondents, 12 (75%) said they could not recall any incident. Three organizations (19%) said they assumed this might have happened at some point, but they could not actually remember a specific incident.

Only a single respondent answered they did recall an 'incident'. When pressed for more details, it turned out that this was a situation where they had not received a security advisory from the National CERT for a vulnerability that was relevant to them, because the organization had not correctly registered a specific asset with the CERT and thus was not receiving advisories pertaining to said asset. This missed advisory did not lead to an actual breach of the asset. The respondent characterized it as a failure of information acquisition, not a security breach.

Interestingly, the three organizations that assumed they probably did miss vulnerability information are all large organizations. They have a relatively high capacity and are processing multiple sources of information – i.e., they are on the left side of our model (Figure 1). The organization that provided the example of a missed vulnerability (which they called an "incident") is a small organization relying mostly on CERT advisories. This suggests a paradox: increased capacity makes organizations more aware of how much vulnerability information is available and, thus, how likely it is that information will be missed.

In sum, no respondent could indicate an incident – in the sense of a security breach – associated with missing vulnerability information. We complemented this finding by asking staff at the National CERT, the sender of security advisories, if they could remember any instance of a breach occurring at an organization related to a vulnerability that the organization had not acquired the information for. They could recall no such instance, while indicating that an incident of this nature would have been very memorable. This means that the organizations experienced no downside or penalty for their limited vulnerability information intake. In the absence of negative feedback from their threat environment, they had no reason to question their coping strategies and limited intake, even though it was mostly driven by their organizational limitations. The limitations in capacity reflect underlying resource allocation decisions, so they are not random, but they are also not the result of formal risk evaluation. Rather they reflect an organic process of trial and error, where the 'error', namely a breach, has not (yet) materialized, as far as they know. Hence, organizations stick with the trial; their initial allocation of (very) limited resources to acquire and process vulnerability information.

When asked if they had changed their sources over the last two years, the majority of respondents focused on the automation of current sources and adding threat intelligence sources, mainly a threat intelligence service provided by the National CERT. Only two respondents stated they had not changed their sources the last two years. Also, only three respondents, all large organizations, described that changes to their sources were part of a structured periodic evaluation. For 12 respondents (75%) the review and changed of sources is not a structured activity but often instigated by external events, such as the new threat intelligence service from the National CERT, or incidents. As one respondent stated: *Changes in our sources are limited. We tried to create an extra source of information by better structuring our threat intelligence.* (9) For two organizations, their changed sources were a result of a new contract for their outsourced SOC services. One respondent was triggered by the question to perform a review, and two others reported that incidents at other organizations prompted them to do this.

9. Discussion

How do organizations deal with the conflict between limited staff time and staying informed about their known vulnerabilities in their attack surface? Of the 22 organizations we interviewed in critical infrastructures and gov-

ernment, not a single one acquired aggregated information via services like NVD or VulnDB, let alone attempt to collect even more comprehensive information. In short, no one drinks from the firehose of *all* published vulnerability disclosures.

The sources of vulnerability information that the participating organizations reported are in line with Li et al. [12]. This prior work did not assess the impact of these sources on the comprehensiveness of the vulnerability information that organizations acquire. We tried to fill this missing gap and contextualized these sources as part of coping strategies to reduce the intake of vulnerability information, in some cases by more than 95%.

9.1. Coping Strategies

The coping strategies reflect organizational constraints. We found that smaller organizations, with smaller IT departments and incident response teams, reduce their intake of vulnerability information much more dramatically than large organizations, who can process the results of automated vulnerability scanning services, covering thousands of new vulnerabilities each year. The coping strategies place varying degrees of trust in the authorities and third parties that provide the sources. The implicit assumption is that these parties will inform the organization about all relevant known vulnerabilities. Some indicated that they assumed the national CERT would notify them about a vulnerability if it was really important, even though the national CERT itself explicitly states that its advisories are not meant to replace independent acquisition of vulnerability information.

Relying on others to learn about about all relevant software vulnerabilities is a strong assumption – and potentially fatal. It helps us to understand why organizations often miss known vulnerabilities. After a breach, it might look negligent or incompetent not to have triaged a vulnerability that was known for months, but yet, in the context of these coping strategies, it is obvious that some vulnerabilities will fall through the cracks, simply because the capacity is lacking for comprehensive triage. In essence, these organizations trust that what comes out of their sources' funnel of information reduction are indeed the most relevant and urgent vulnerabilities for them. The basis for this trust is not really clear, but, as we argued in Section 8, it is a rational strategy to follow. None of the interviewees recalled their organizations experiencing an incident associated with vulnerability information that they had missed (although some speculated that one might have occurred in the past). There is no feedback loop that suggests that their trust in, say, the national CERT advisories, was not warranted.

One example of relying on others is to use threat intelligence providers as a source of vulnerability information. While threat intelligence covers only a fraction of all vulnerabilities, it informs organizations about vulnerabilities that are actively being exploited. Given that most vulnerabilities are never exploited in the wild [31, 32], it makes sense to focus scarce resources on those cases. In the U.S., this practice is now formally codified. The U.S. Cybersecurity & Infrastructure Security Agency (CISA) publishes a Known Exploited Vulnerabilities Catalog [37] and it issued a Binding Operational Directive [38] that requires all federal civilian agencies to patch within two weeks all vulnerabilities added to this catalog. CISA also encourages all private organizations and local governments to do so.

Still, the amount of trust placed, implicitly or explicitly, in the entities that provide vulnerability information raises questions about responsibility and accountability. Who is held accountable when organizations only learn about a vulnerability after an incident? Do their providers of vulnerability information have any responsibility? For example: can breached organizations blame CISA for not including the exploited vulnerability in their catalog, even though it was clearly attacked in the wild? Does the national CERT share some of the responsibility by sending out more than a thousand advisories every year? Do commercial providers of scanning services face any liability for not including a vulnerability or not adequately detecting it? Or does all responsibility fall on the organization itself? As long as organizations do not experience security breaches because of missed vulnerability information, the answers to these questions remain mostly theoretical.

Also, we have to recognize that it is already hard enough to act on the limited vulnerability information that the organizations currently acquire. It is one thing to learn about a software vulnerability, getting that vulnerability patched is something else altogether. It requires overcoming organizational inertia, often because of the business continuity impacts of patching. Respondents mentioned that one reason they heavily rely on the advisories of the national CERT, is that these advisories carry the authority needed to convince their own organizations to initiate a patching process. Their management assigns more priority to these security advisories than to other sources of vulnerability information. In those circumstances, it might not make sense to acquire more information about known vulnerabilities, if it is already hard enough to remediate the ones you do know about.

9.2. Recommendations

So what recommendations can we draw from these findings? First of all, our findings call into question the conventional security advice that organizations should try to acquire comprehensive vulnerability information for their assets, in order not to miss any known vulnerabilities [33, 34]. The organizations we studied do not follow this advice but are not suffering negative consequences for their coping strategies. Of course, one could argue that the lack of experienced breaches only means they have been lucky or perhaps even ignorant, if they were breached but did discover it. We cannot rule this out. There is always the possibility of low probability/high consequence events, i.e., 'black swans', meaning that the strategies of our respondents are not immune to failure. That being said, it would be wrong to discount the learning experience of these organizations: their limited intake of vulnerability information has proved good enough, so far.

Acquiring and processing more vulnerability information is not free. This is clearly shown by the association we find between organizational capacity and the magnitude of vulnerability information being consumed. So these organizations face a trade-off: should they divert resources away from elsewhere, perhaps from the patching process itself, to take in more vulnerability information? How would they know that this would imply a net gain in terms of security? We would argue that we cannot know that. The lack of incidents suggest there is little evidence that more information will increase security and the diverted resources might reduce security efforts elsewhere. In short, the conventional wisdom that acquiring more comprehensive information is better, is not supported by the evidence. Our recommendation is to re-evaluate this advice and view the coping strategies we uncovered as rational (in the sense of bounded rationality [39]) and as a legitimate strategy.

Our second recommendation builds on the finding that these practices have, for the most part, grown organically. It is rarely the outcome of an explicit evaluation process. Organizations seemed mostly unaware of how much information they are missing. They left the evaluation of their information intake as an afterthought. This may be because better curation is always important, but never urgent. We propose that organizations periodically review their intake of vulnerability information, for example on a yearly basis, or whenever service renewal comes up. Smaller organizations can suffice with an informal discussion among team members. Larger organizations should consider formalizing this discussion with suitable metrics and measurements – because as the attack surface grows, so does the likelihood of experiencing an incident due to incomplete vulnerability information. We suggest asking the following questions: *i)* Which sources of vulnerability information cover which of our assets? *ii)* Do we have assets that are currently not tracked by any source? *iii)* Can we trust the prioritization done by our sources? *iv)* Are our sources timely and do they contain actionable mitigation information?

Finally, we suggest to pay attention to vulnerability information in post-mortem of security incidents in which software vulnerabilities were exploited. Three respondents reported that they assumed their organization must have had incidents associated with missing information, even though they could not remember an actual example. Incidents provide an important signal to evaluate and improve sourcing of vulnerability information, in terms of checking whether the organization had received and processed information related to the exploited vulnerability. Going forward, a more formalized foundation is to include the vulnerability information sources in the risk analysis process. This would then also allow coping strategies such as those that we identified to be explicitly chosen, rather than an implicit outcome.

10. Limitations

The design of our study implies four key limitations. First of all, the sample of respondents and their organizations was recruited within one country (the Netherlands) and via its national CERT. All participants came from critical infrastructures and government services. While the patterns that we found fit well within prior work conducted in other contexts, we have to be careful in generalizing these findings to other organizations and other countries.

Second, we interviewed only one respondent per organization. These interviewees had first-hand knowledge of the practices we wanted to observe, since they are the actual recipients of security advisories and other sources of vulnerability information. That said, their knowledge about the practices in other parts of the organization is likely to be incomplete. It seems unlikely that wholly separate vulnerability triage processes are going on that they are unaware of, it is possible that elsewhere in the organization people also acquire some vulnerability information, especially in the larger organizations.

Third, we tried to quantify the volume of vulnerability information that organizations ingest, but these sources are highly variable across time and providers. Also, respondents might not accurately remember the actual volumes of certain flows. Our attempts at quantification are meant primarily as a way to qualitatively compare the sources, not to pin them down to actual numbers. We have a reasonable idea of the order of magnitude of the total pool of discovered vulnerabilities, based on sources like VulnDB. We ranked the other sources in terms of their order of magnitude compared to that initial set and to each other.

Last, since we recruited participants via the national CERT, it is possible that their answers reflect a social desirability bias towards questions on the CERT advisories. We tried to counter potential bias by asking indirect questions and rephrasing questions to cover all forms of vulnerability information.

11. Conclusion

A range of catastrophic breaches all share as their root cause an unpatched, but known, vulnerability. We identified coping strategies that organizations use to acquire, but also reduce, the intake of software vulnerability information. The lack of comprehensive acquisition of knowledge about published vulnerabilities is a disconcerting finding. Yet, it can also be understood from the resource constraints that organizations face. These organizations could not remember suffering any incident associated with missed vulnerability information. While there is no guarantee that there was no undiscovered breach or that no breach will occur in the future, the organizations' experience suggest that it might be rational to be very selective in the intake of information, contrary to accepted best practices. Increasing that intake has no visible security benefits to them, as most vulnerabilities are never attacked and therefore no breaches would necessarily be prevented, while it does require resources that might be more effective elsewhere. We do recommend that organizations conduct more deliberate evaluations of their vulnerability information acquisition, assessing the often implicit trade-offs in a more formal risk management process.

Acknowledgments

The authors are grateful to Matthijs Krijgsman for his help with coding the many interviews. The authors would also like to thank Jonathan Spring at CMU and the anonymous reviewers for their comments that have been helpful in improving the paper. This research has been partially funded by the Ministry of the Interior and Kingdom Relations of the Netherlands and partially by the Dutch Research Council (NWO) as part of the THESEUS project (NWA.1215.18.006).

References

[1] Susan Moore. *Focus on the Biggest Security Threats, Not the Most Publicized.* Nov. 2018. URL: https://www.gartner.com/smarterwithgartner/focus-on-the-biggest-security-threats-not-the-most-publicized/.

[2] Dan Goodin. *Failure to patch two-month-old bug led to massive Equifax breach*. Sept. 2017. URL: https://arstechnica.com/information-technology/2017/09/massive-equifax-breach-caused-by-failure-to-patch-two-month-old-bug/.

[3] Andy Greenberg. *The Untold Story of NotPetya, the Most Devastating Cyberattack in History*. Aug. 2018. URL: https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/.

[4] Brian Carlson. *The Microsoft Exchange Server hack: A timeline*. May 2021. URL: https://www.csoonline.com/article/3616699/the-microsoft-exchange-server-hack-a-timeline.html.

[5] Forum of Incident Response and Security Teams (FIRST). *CSIRT Services Framework 2.1*. Aug. 2021. URL: https://www.first.org/standards/frameworks/csirts/csirt_services_framework_v2.1.

[6] Noura Alomar et al. ""You've Got Your Nice List of Bugs, Now What?" Vulnerability Discovery and Management Processes in the Wild". In: *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. USENIX Association, Aug. 2020, pp. 319–339. ISBN: 978-1-939133-16-8. URL: https://www.usenix.org/conference/soups2020/presentation/alomar.

[7] *VulnDB Vulnerability Statistics*. URL: https://vulndb.cyberriskanalytics.com/#statistics.

[8] Lucas Miranda et al. "On the Flow of Software Security Advisories". In: *IEEE Transactions on Network and Service Management* (2021), pp. 1–1. DOI: 10.1109/TNSM.2021.3078727.

[9] Stefan Frei et al. "Large-scale vulnerability analysis". In: *Proceedings of the 2006 SIGCOMM workshop on Large-scale attack defense*. 2006, pp. 131–138.

[10] Carl Sabottke, Octavian Suciu, and Tudor Dumitras. "Vulnerability Disclosure in the Age of Social Media: Exploiting Twitter for Predicting Real-World Exploits". In: *24th USENIX Security Symposium (USENIX Security 15)*. Washington, D.C.: USENIX Association, Aug. 2015, pp. 1041–1056. ISBN: 978-1-939133-11-3. URL: https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/sabottke.

[11] Su Zhang, Xinming Ou, and Doina Caragea. "Predicting cyber risks through national vulnerability database". In: *Information Security Journal: A Global Perspective* 24.4-6 (2015), pp. 194–206.

[12] Frank Li et al. "Keepers of the Machines: Examining How System Administrators Manage Software Updates For Multiple Machines". In: *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. Santa Clara, CA: USENIX Association, Aug. 2019. URL: https://www.usenix.org/conference/soups2019/presentation/li.

[13] Jessica Staddon and Noelle Easterday. ""it's a generally exhausting field" A Large-Scale Study of Security Incident Management Workflows and Pain Points". In: *2019 17th International Conference on Privacy, Security and Trust (PST)*. IEEE. 2019, pp. 1–12.

[14] Constanze Dietrich et al. "Investigating system operators' perspective on security misconfigurations". In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 2018, pp. 1272–1289.

[15] Platon Kotzias et al. "Mind Your Own Business: A Longitudinal Study of Threats and Vulnerabilities in Enterprises." In: *Network and Distributed System Security Symposium (NDSS) 2019*. 2019.

[16] F. Li and V. Paxson. "A Large-Scale Empirical Study of Security Patches". In: *CCS '17: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. New York, NY, USA: Association for Computing Machinery, Oct. 2017, pp. 2201–2215. ISBN: 978-1-45034946-8. DOI: 10.1145/3133956.3134072.

[17] *CVSS v3.1 Specification Document*. URL: https://www.first.org/cvss/specification-document.

[18] Jonathan Spring et al. "Time to Change the CVSS?" In: *IEEE Security Privacy* 19.2 (2021), pp. 74–78. DOI: 10.1109/MSEC.2020.3044475.

[19] Luca Allodi and Fabio Massacci. "Comparing vulnerability severity and exploits using case-control studies". In: *ACM Transactions on Information and System Security* 17.1 (2014). ISSN: 15577406. DOI: 10.1145/2630069.

[20] Jonathan Spring et al. *Towards Improving CVSS*. 2018. URL: https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=538368.

[21] Luca Allodi et al. "The Effect of Security Education and Expertise on Security Assessments: the Case of Software Vulnerabilities". In: *arXiv* (Aug. 2018), pp. 1–15. ISSN: 23318422. arXiv: 1808.06547. URL: http://arxiv.org/abs/1808.06547.

[22] Fengli Zhang and Qinghua Li. "Dynamic Risk-Aware Patch Scheduling". In: *2020 IEEE Conference on Communications and Network Security (CNS)*. 2020, pp. 1–9. DOI: 10.1109/CNS48642.2020.9162225.

[23] K. Alperin et al. *Risk Prioritization by Leveraging Latent Vulnerability Features in a Contested Environment*. New York, NY, USA: Association for Computing Machinery, Nov. 2019. ISBN: 978-1-45036833-9. DOI: 10.1145/3338501.3357365.

[24] Ali Alshawish and Hermann De Meer. "Prioritize When Patching Everything is Impossible!" In: *Proceedings - Conference on Local Computer Networks, LCN* 2019-October (2019), pp. 125–128. DOI: 10.1109/LCN44214.2019.8990847.

[25] Jay Jacobs et al. "Improving vulnerability remediation through better exploit prediction". In: *Journal of Cybersecurity* 6.1 (Sept. 2020). ISSN: 2057-2085. DOI: 10.1093/cybsec/tyaa015. eprint: https://academic.oup.com/cybersecurity/article-pdf/6/1/tyaa015/33746021/tyaa015.pdf. URL: https://doi.org/10.1093/cybsec/tyaa015.

[26] Chaowei Xiao et al. "From Patching Delays to Infection Symptoms: Using Risk Profiles for an Early Discovery of Vulnerabilities Exploited in the Wild". In: *27th USENIX Security Symposium (USENIX Security 18)*. Baltimore, MD: USENIX Association, Aug. 2018, pp. 903–918. ISBN: 978-1-939133-04-5. URL: https://www.usenix.org/conference/usenixsecurity18/presentation/xiao.

[27] Xander Bouwman et al. "A different cup of TI? The added value of commercial threat intelligence". In: *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020, pp. 433–450. ISBN: 978-1-939133-17-5. URL: https://www.usenix.org/conference/usenixsecurity20/presentation/bouwman.

[28] Virginia Braun and Victoria Clarke. "Using thematic analysis in psychology". In: *Qualitative research in psychology* 3.2 (2006), pp. 77–101.

[29] Susanne Friese. *Qualitative data analysis with ATLAS.TI*. Sage, 2019.

[30] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. "Reliability and Inter-Rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice". In: *Proc. ACM Hum.-Comput. Interact.* 3.CSCW (Nov. 2019). DOI: 10.1145/3359174. URL: https://doi.org/10.1145/3359174.

[31] Allen D Householder et al. "Historical analysis of exploit availability timelines". In: *13th USENIX Workshop on Cyber Security Experimentation and Test (CSET 20)*. 2020.

[32] Luca Allodi and Fabio Massacci. "A preliminary analysis of vulnerability scores for attacks in wild". In: *Proceedings of the 2012 ACM Workshop on Building analysis datasets and gathering experience returns for security - BADGERS '12*. New York, New York, USA: ACM Press, 2012, p. 17. ISBN: 9781450316613. DOI: 10.1145/2382416.2382427. URL: http://dl.acm.org/citation.cfm?doid=2382416.2382427.

[33] Carnegie Mellon. *CRR Supplemental Resource Guide: Volume 4 Vulnerability Management*. Apr. 2016. URL: https://www.cisa.gov/sites/default/files/publications/CRR_Resource_Guide-VM_0.pdf.

[34] NIST. *Framework for Improving Critical Infrastructure Cybersecurity*. Apr. 2018. URL: https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf.

[35] Irfahn Khimj. *The Five Stages of Vulnerability Management Maturity*. Apr. 2019. URL: https://www.tripwire.com/-/media/tripwiredotcom/files/white-paper/tripwire_the_five_stages_of_vm_maturity_white_paper.pdf.

[36] Jonathan Risto. *The Five Stages of Vulnerability Management Maturity*. Oct. 2020. URL: https://www.sans.org/blog/vulnerability-management-maturity-model-part-ii.

[37] U.S. Cybersecurity and Infrastructure Security Agency (CISA). *Known Exploited Vulnerabilities Catalog*. Nov. 2021. URL: https://www.cisa.gov/known-exploited-vulnerabilities-catalog.

[38] U.S. Cybersecurity and Infrastructure Security Agency (CISA). *Binding Operational Directive 22-01*. Nov. 2021. URL: https://www.cisa.gov/binding-operational-directive-22-01.

[39] H. A. Simon. "Bounded Rationality". In: *Utility and Probability*. London, England, UK: Palgrave Macmillan, 1990, pp. 15–18. ISBN: 978-0-333-49541-4. DOI: 10.1007/978-1-349-20568-4_5.

Appendix

1. Survey questions

Below the questions are listed from our survey, conducted after the interview, as stated in section 3

1) What is your job title?
   [OpenAnswer]
2) How many years are you employed in your current job function?
   [Number]
3) Where do you work?
   SOC
   CERT/CSIRT
   Something else, namely ...
4) How large is your department in terms of staff?
   Less than 5
   5 – 10
   11-25
   26-50
   51-100
   101- 200
   More than 200
5) How large is your incident response in terms of staff?
   Less than 5
   5 – 10
   11-25
   26-50
   51-100
   101- 200
   More than 200
6) How many employees does your organization have?
   Less than 50
   51 – 100
   101-500
   501-2000
   2001 – 5000
   5001- 10.000
   More than 10.000
7) In which industry/sector do you work?
   [OpenAnswer]

2. Interview protocol

We conducted our interviews, as described in section 3 along the following set of questions:

The inflow of vulnerability information

1) In what mailbox do you receive [REDACTED] security advisories?
2) Do you receive all or a selection of these advisories?
3) If so, what is this selection based on?
4) How would you describe this mailbox?
5) How intensely is this mailbox monitored?
6) How many individuals monitor this mailbox simultaneously?
7) What is the job function of these individuals?
8) What is your opinion on receiving the information in this manner?
9) In terms of amount of information, what are the five largest sources of vulnerability information you receive?
10) How many instances of vulnerability information do you receive in a day?
11) Which of these sources are paid sources?
12) Do you receive these sources in different mailboxes?
13) If so, what is the ratio of this division?
14) How intensely is this information monitored?
15) How many individuals monitor this mailbox simultaneously?
16) What is the job function of these individuals?
17) On what other communication channels do you receive vulnerability information?
18) How do you receive information on critical vulnerabilities?

The vulnerability triage process

1) Assessment of vulnerability information
2) How do you process vulnerability information?
3) What source do you read first and why?
4) Who assesses the risk of a potential vulnerability?
5) What do you do with vulnerability information about potential critical vulnerabilities?
6) How would you rank these types of vulnerability information sources in order of importance and why?
7) How much overlap is between these sources?
8) What information in security advisories is most/least important?
9) Do you use the information in security advisories and how?
10) Processing vulnerability information
11) What departments in your organization receive notifications of vulnerabilities?
12) How are vulnerability information sources used by the rest of the organization?
13) What source receives the most attention by the rest of the organization and why?
14) How are you involved in the rest of the vulnerability management process?
15) What would you like improved in your vulnerability management process?

3. Follow-up interview protocol

1) Were your organization's sources of vulnerability information actively selected or are they rather the result of an organic process?
2) Are you under the impression that you may be missing possibly relevant information?
3) Has your organization experienced an incident or situation around a vulnerability that your sources had missed?
   a) If affirmative: Did this lead to adjustments in your use of sources of vulnerability information?
   b) If negative: Has your organization over the last two years made changes to your use of sources of software vulnerability information, and why?

TABLE 6. KEY CODES PER TRADE-OFF VARIABLE

Description	*Key codes*
Independently assessing vulnerability information in-house	Lean forward; CVE; independent assessment; External Collaboration; Triage proces; independent information source; Urgency in critical vulnerability; Formal process; Trust external parties in vulnerabilities; Importance of Vendor information; Clustering of information sources;
Trusting authoritative peers or third parties to assess and filter information	Trust [REDACTED]; Trust external parties; reliance external parties; Not enough capacity; External risk assessment determines triage; ad hoc information; lean back solutions; ease; collaboration externally; trust vendors; trust external parties supplying information; external SOC; collaboration externally; type of vulnerability leading for collaboration; small organization; CVE after advisory
Proactively search and assess information	Active monitoring; Assets known; Afraid to miss vulnerabilities; Priority status; Own assessment determines urgency; treatment critical vulnerability; seperate risk assesment SOC; Important role scanners information management; Continuous informing of management; automated receiving vulnerability information; lean forward with urgency; vulnerability list for management; searches proactively for CVE
Reactive activities on information	Lean back asset management; No paid other sources; lacking sight; happy other sources confirm critical vulnerability; No active search for vulnerabilities, only passive from external sources; trust in vendors; Assessment [REDACTED] important for actions management; Waiting monthly report remediate vulnerability; [REDACTED] assessment forces action; wait and see; filtering high vulnerabilities; only high vulnerabilities own risk assessment
Formalized processing of information	Active monitoring; Dev Ops/Agile; formalized processes; decision management critical vulnerability; ticketing system; separation of roles within process; periodic monitoring; First handler in triage process; high capacity because of automation;
Ad hoc processing of information	Ad hoc information inflow; small team personal management; monitoring sources out of personal interest; public sources; Manual assessment; Lean forward; multiple roles within vulnerability process; role of system admin; Unstructured information sources; Limited capacity, desire to automate; Personal networks; Individual responsible for solution