The Unpatchables
Why Municipalities Persist in Running Vulnerable Hosts

Ethembabaoglu, Aksel; van Wegberg, Rolf; Zhauniarovich, Yury; van Eeten, Michel

Publication date
2024

Document Version
Final published version

Published in
Proceedings of the 33rd USENIX Security Symposium

Citation (APA)
Ethembabaoglu, A., van Wegberg, R., Zhauniarovich, Y., & van Eeten, M. (2024). The Unpatchables: Why Municipalities Persist in Running Vulnerable Hosts. In *Proceedings of the 33rd USENIX Security Symposium* (pp. 7049-7066). USENIX Association.
https://www.usenix.org/conference/usenixsecurity24/presentation/ethembabaoglu

The Unpatchables: Why Municipalities Persist in Running Vulnerable Hosts

Aksel Ethembabaoglu, Rolf van Wegberg, Yury Zhauniarovich,
and Michel van Eeten, *Delft University of Technology*

https://www.usenix.org/conference/usenixsecurity24/presentation/ethembabaoglu

This paper is included in the Proceedings of the
33rd USENIX Security Symposium.

August 14–16, 2024 • Philadelphia, PA, USA

978-1-939133-44-1

The Unpatchables: Why Municipalities Persist in Running Vulnerable Hosts

Aksel Ethembabaoglu, Rolf van Wegberg, Yury Zhauniarovich, and Michel van Eeten

Delft University of Technology

Abstract

Many organizations continue to expose vulnerable systems for which patches exist, opening themselves up for cyberattacks. Local governments are found to be especially affected by this problem. Why are these systems not patched? Prior work relied on vulnerability scanning to observe unpatched systems, notification studies on remediating them, and on user studies of sysadmins to describe self-reported patching behavior, but they are rarely used together as we do in this study. We analyze scan data following standard industry practices and detect unpatched hosts across the set of 322 Dutch municipalities. Our first question is: Are these detections false positives? We engage with 29 security professionals working for 54 municipalities to collect ground truth.

All detections were accurate. Our approach also uncovers a major misalignment between systems that the responsible CERT attributes to the municipalities and the systems the practitioners at municipalities believe they are responsible for. We then interviewed the professionals as to why these vulnerable systems were still exposed. We identify four explanations for non-patching: *unaware*, *unable*, *retired* and *shut down*. The institutional framework to mitigate cyber threats assumes that vulnerable systems are first correctly identified, then correctly attributed and notified, and finally correctly mitigated. Our findings illustrate that the first assumption is correct, the second one is not and the third one is more complicated in practice. We end with reflections on how to better remediate vulnerable hosts.

1 Introduction

Exploiting known vulnerabilities for which a patch exists remains a dominant attack vector, even after years of warnings [11]. Local governments are seen as especially susceptible [26, 37]. The U.S. Cybersecurity and Infrastructure Security Agency (CISA) referred to them as the "cyber poor", offering vulnerability scanning services as support [31]. These concerns are not exclusive to the U.S. In the Netherlands, the Dutch Safety Board investigated the incidents following the 2020 Citrix vulnerabilities and concluded that municipalities struggle with patching because of a lack of resources [17].

The threat of exploitation of local governments, or any other organization, is not hypothetical. Municipalities worldwide have been hit with ransomware attacks paralyzing organizations and losing sensitive and personal information of citizens [10, 24, 43, 45]. These attacks had destabilizing societal effects, with governmental services being unavailable and data of citizens being lost. In the US alone, more than a hundred local government organizations reported cyberattacks in 2019 and 2020 [42].

To mitigate such threats, governments established Computer Emergency Response Teams (CERTs) or Computer Security Incident Response Teams (CSIRTs) [19]. These organizations receive security incident data and network scan information from various sources. This data is forwarded to the organization responsible for the vulnerable systems. The notified organization is then expected to mitigate the vulnerability. Prior research shows that such security notifications can expedite vulnerability remediation [32, 52]. CERTs around the world operate on a similar model of monitoring networks and notifying constituents. In Brazil, the CERT provides incident analysis and coordination services for any network that uses IP addresses or Autonomous Systems allocated to Brazil, and domains under the .br ccTLD. It alerts Brazilian networks involved in malicious activities [8]. The CERT-Bund in Germany supports handling IT security incidents; it provides active alerts for the federal administration in the event of acute threats [21]. In Africa, the non-profit organization AfricaCERT, with several African countries as its members, states in its objectives that it "encourages information sharing in ICT security, which includes findings from reported incidents and case studies, so that vulnerabilities can be rapidly identified, and its risks mitigated" [2]. In the UK, the NCSCS introduced the Early Warning service, offering its members to notify them of vulnerabilities in their networks [38]. These CERTs monitor threats and attribute IP networks to organizations for alerting and notifications.

Yet, despite these measures, governmental systems, as well as other organizations, are still frequently found to be vulnerable [49]. Therefore, in this study, we first ask the question: how accurate are the measurements of unpatched systems? Next, we consider the question: why are those systems not patched? Prior work has shed light on the presence of unpatched systems via three strands of research, (*i*) studying the self-reported patching practices of system administrators [13, 33], (*ii*) network scans that detected where vulnerable systems are located [27, 56], and (*iii*) the effectiveness of notifications about detected vulnerable systems [32, 52]. However, no prior work has integrated these three strands as we do in this study.

We combine passive network scans (relying on banner-grabbing to infer software versions) to link versioned software to known vulnerabilities from the National Vulnerability Database (NVD) and to detect vulnerable systems in local governments in the Netherlands that receive notifications from their CERT and then use these detections in interviews with security practitioners responsible for those systems. We gather ground truth on the detections to answer our first research question. Next, to answer our second question, we explore the non-patching behavior of practitioners in a way less sensitive to the biases that come with self-reporting, which previous studies have relied on, e.g., [33].

We collaborated with IBD-CERT, the CERT organization for all municipalities in the Netherlands. The municipalities have registered their IP network ranges with the IBD-CERT. The CERT receives scan data from the national CERT and other sources about hosts with CVEs (Common Vulnerabilities and Exposures). It then notifies its members about detected hosts in their networks. We explore potential explanations for the detected vulnerable systems. First, the passive network scans might produce false positives: the host might not actually be running the vulnerable software. Scan data can contain artefacts and version information from hosts that are manipulated or simply wrong. Second, the systems are unpatched, but there is a reason the municipality has not patched it. It might be unaware of the vulnerability, it might be unable to patch it, or it might have decided that patching is not needed. In the process of conducting the interviews, a third explanation arose: the municipalities do not consider the vulnerable system to be their responsibility, even though they reside in the IP ranges that they registered with the CERT.

We analyzed 1,687 registered IP ranges covering 322 municipalities in the country. Using passive scanning data from Shodan [35] and Censys [14], we observed 154 vulnerable hosts running 17 different services with 643 unique CVEs in total. We conducted 16 semi-structured interviews with 29 security practitioners working for 54 municipalities (some IT departments support multiple municipalities), covering about 17% of the total population of Dutch municipalities. This sample includes municipalities with and without exposed vulnerabilities, so we can compare their answers and the features of their organizations. We transcribed, coded, and analyzed the interview data and conducted follow-up conversations.

First, we observed that the observation of vulnerable hosts seems reliable and not plagued by false positives. Next, we found that a significant portion of the vulnerable systems that get notified about fall into a gap, because there is a misalignment between the IPs of municipalities registered at the CERT, and the IPs the sysadmins see themselves as responsible for. At least some of these vulnerable systems appear to be "shadow IT". This might explain why many of these systems persist in a vulnerable state. It also explains that the municipalities see themselves as much less vulnerable than their CERT or central government does. For the systems that were administered by the municipality, we observed that respondents were (*i*) not aware of vulnerable hosts, (*ii*) unable to patch the system, or (*iii*) systems were in the process of being retired. We also learned that vulnerable systems are rarely shut down because of security reasons. We make the following contributions:

- We collect ground-truth evidence that the external detection of vulnerable services is accurate and not plagued by false positives. The observations we collected from Shodan and Censys appeared to be 100% correct.
- We demonstrate major misalignments between the systems the CERT attributes to a municipality and the systems a municipality believes it is responsible for. For our sample of municipalities, this misalignment translates to the CERT observing 18 vulnerable hosts that the IT departments do not consider their responsibility, pointing to the problem of "shadow IT". On the other hand, the municipal IT departments do see themselves as responsible for 6 vulnerable hosts that the CERT does not attribute to them and thus doesn't notify them about. Only 9 vulnerable hosts are seen and attributed consistently by both organizations. These observations raise concerns about the effectiveness of the incident response framework of CERTs for notifying victim organizations.
- We identify four categories for not patching vulnerable systems in practice: *unaware*, *unable*, *retired*, and *shutdown*. In most cases, security professionals were unaware of the vulnerable system. Additionally, we find that there are no CVE or application-specific mitigation strategies applied to vulnerable hosts unless explicitly provided in the security advisory of the CERT or from the vendor. We also observe a strong tension between business continuity and security in the vulnerability management process.

2 Related Work

Our study ties into three main strands of research: (*i*) user studies of security practitioners or IT professionals; (*ii*) studies using network scans to collect observational data on vulnerable systems; and (*iii*) studies on security and vulnerability notifications. We discuss each in turn.

First, several studies examined the perspectives of security practitioners to gain an understanding of their considerations or the organizational processes in which they operate, not necessarily related to security [5, 55]. Li et al. [33] identified processes system administrators use to manage software updates but relied solely on self-reporting via surveys. The software updates that system administrators reported were not empirically measured, and therefore a picture may be painted that does not fully align with reality. Dietrich et al. [13] looked at how system administrators managed their systems and their configurations, specifically examining the perspectives of system administrators. Velasquez et al. [53] looked at the role of the system administrator within the organization and found that they often act as a broker between the end-users and the technical community. Krombholz et al. [28] found that the deployment process of security measures for system administrators is too complex and recommended that server configurations should opt for security by default. Alomar et al. [3] observed that practitioners struggle with vulnerability remediation and that vulnerability discovery efforts are hindered by significant trust, communication, funding, and staffing issues. Smale et al. [12] found that vulnerability information acquisition by practitioners is not comprehensive and that up to 95% of all CVE disclosures are not ingested. These studies examine the perspectives of practitioners, but they are all based on self-reported behaviors, which is potentially biased. By using actual network data linked to known vulnerabilities, we ground the responses of the interviewees by discussing with them the evidence of vulnerable hosts in their network.

Vulnerabilities can be found in the wild with passive scanning services. The work of O'Hare et al. [44] presents a method to discover vulnerabilities by combining the CPE from passive scanning services with data from the National Vulnerability Database (NVD). A vast area of research relates to the use of Internet-wide scans with Zmap [16] and similar tools to detect vulnerable hosts [18, 23]. Numerous studies used Censys and Shodan to measure vulnerable systems [6, 15]. The work by West et al. [56] found that the Internet-facing OpenSSH service might not be as vulnerable as initially suspected due to the use of backports. Kotzias et. al [27] observed that the patching of server applications is much slower than the patching of client-side applications.

Finally, as stated in the introduction, there exists a line of research related to incident response and victim notification [52]. Li et al. observed that vulnerability notifications addressed directly to the owners of the resources promoted faster remediation than those sent to national CERTs [32]. Cetin et al. showed that retrieving contact information at scale was problematic. But once contacted, entities were more likely to remediate [9].

Our work builds on previous studies by connecting and contextualizing different methods and data sources to provide a deeper understanding of patching behavior and the responsibilities of networks. The combination of scanning networks and using that data in interviews should mitigate the risk of self-reporting in patching behavior. Additionally, it allows us to verify external measurements of software versions of networks, to obtain ground-truth. We complement those external network measurements with qualitative data to record the considerations of practitioners as to why those vulnerable systems exist in their infrastructure. Lastly, by collaborating with the CERT and the municipalities we are able to correlate (assumed) responsibilities for specific IP addresses that are attributed the municipalities.

3 Ethics

We received approval from our Institutional Review Board for conducting this human-subjects research. Participants were explained in detail about the study, associated risks, and use of information for which they provided informed consent.

Research on the vulnerabilities of an organization is a sensitive topic. In our informed consent form, as well as in the recruitment emails for the interviews, we consistently assured the participating municipalities, as well as their CERT, that their data was handled confidentially and would only be presented in an aggregated and anonymized form. No identities or municipalities would be named. We also made sure that answers that referred to specific tooling that might reveal their identity were cleaned.

After the interviews, we reported IPs with vulnerable services to the CERT. During the interviews, we notified respondents of the observed vulnerable hosts.

To minimize the burden on the security staff, we choose to use passive rather than active scans in order to prevent the disruption of regular operations or unintentionally triggering alerts (false positives) in their Security Operation Centers (SOCs). Before publication, we presented a draft of this paper to the respondents, so that they had a chance to check and correct quotes attributed to them.

4 Measurement Approach

We used a mixed-methods approach that combines external network measurements of the municipalities with a qualitative user study among sysadmins and security practitioners. The aim of this approach is two-fold. First, we validate the passive network measurements for detecting software versions with the responsible operators. This allows us to estimate the false positive ratio for the detection of vulnerable systems when relying on banner information about software versions and linking those to known vulnerabilities. We did not carry out active measurements on the networks of municipalities to verify vulnerabilities. Second, we conducted interviews to learn from the practitioners why the vulnerable systems, assuming they were correctly detected, were present in their network. We collaborated with the IBD-CERT, which is responsible for all Dutch municipalities. The CERT provided us with the IP ranges that the municipalities have registered with it. It also facilitated the process of recruiting interviewees.

4.1 Scanning Municipal Networks

Network scans can be done actively or passively. Active scans directly connect to the target network. They grab banners and infer software versions (e.g., Nmap [34]) but they can also be more intrusive. Some tools such as Metasploit [47], Nessus [51] or Qualys [46] actually try to exploit a potential vulnerability to determine if it is present on the target system.

By contrast, passive scanners run their own Internet-wide banner-grabbing scans and present the results to its users, often as a service via a web portal. The scan data is (somewhat) outdated but the target network is not directly touched by the users of the service. Popular passive scanning services are Shodan [35] and Censys [14].

The CERT provided the research team with 1,687 IP ranges for 322 municipalities. In the Netherlands, there are 346 municipalities, giving us coverage of over 93% of the total population. To reduce the burden on the municipality networks, we relied solely on passive rather than active scans. In November 2022, we queried Censys and Shodan to identify hosts. This process resulted in 3402 detected hosts. Figure 1 depicts the 322 municipalities from the CERT and the number of responding hosts for each. The next step was to determine what services were running on the detected hosts. Censys and Shodan parse banners to determine the service and version that is running on a port. A portion of the service banners contained version information. The service and, if available, version are used to generate a Common Platform Enumeration (CPE) identifier. We collected CPE identifiers of the applications and their versions running on a host. For those CPEs, we retrieved the accompanying Common Vulnerability Disclosures (CVEs) from the National Vulnerability Database (NVD) [40]. For these CVEs, we also collected their Common Vulnerability Scoring System Version 3 (CVSS3) scores. CVSS3 is an open framework for communicating the characteristics and severity of software vulnerabilities [20]. We used the CVSS score to label vulnerabilities as Critical, High, and Medium/Low [41].

Our study primarily relies on the scan results from Censys. Contrary to Shodan, it does not perform any blackbox post-processing, so we can more transparently infer CPEs. Furthermore, Censys runs its scanners on a daily basis for each IP address [7]. To corroborate its results, we compared the results to those we got from Shodan. We observed two minor discrepancies. First, we found that for some IP ranges, each service returned a different number of hosts. (Both services did return the same number of hosts that were running a versioned service.) We asked respondents if they actively blocked either Censys or Shodan, which none did. However, several respondents did mention that their firewall blocks consecutive requests from a scanner and that it might explain the difference in resulting hosts. Second, we found minor discrepancies in how Shodan and Censys parse the banner for the CPE. For example, Shodan detects unversioned Apache2 instances conveying the following CPE for them: *cpe:2.3:a:apache:http_server:2*, while Censys returned the *cpe:2.3:a:apache:http_server:*:*:*:*:*:*:*:** CPE. Also, for Nginx 1.18.0, Shodan returns the CPE *cpe:2.3:a:igor_sysoev:nginx:1.18.0* while Censys returns *cpe:2.3:a:nginx:nginx:1.18.0:*:*:*:*:*:*:**. For neither CPE did it have an effect on the associated vulnerabilities.

4.2 User Study

Selection of municipalities. We compiled a list of versioned and (vulnerable) services per municipality. For the interviews, we selected municipalities on three criteria. First, we preferred municipalities with the highest number of hosts running versioned services, to maximize the number of external measurements of systems. We interviewed 10 municipalities with vulnerable hosts. Second, we interviewed municipalities without vulnerable systems. We wanted to learn if they approached vulnerability management differently, or if their organizational structure may have an effect. We interviewed 6 municipalities without vulnerable hosts. Third, we wanted a diverse set of municipalities in terms of size and geographical location. We aimed for a diverse sample measured in the number of inhabitants and Internet-facing hosts. Depicted in Figure 1, we plotted in red the participating municipalities among the total set of municipalities.

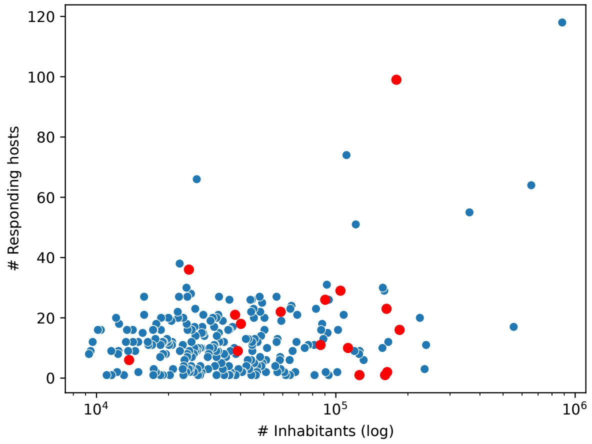

Figure 1: **Responding hosts versus the number of inhabitants (log scale) for the IPs of all municipalities. Red dots are the municipal IT organizations (16) we interviewed.**

Based on these criteria, we selected 34 municipalities and, in collaboration with the CERT, reached out to them in November 2022 until April 2023 via e-mail, inviting them to join a one-hour semi-structured interview. In case of no reply, we sent out a reminder after several weeks. We asked for respondents who were involved in the operational process of detecting and mitigating vulnerabilities, often system administrators. In total, we were able to set up interviews (4 declined to participate, and 12 did not respond). In several cases, municipalities maintained a shared ICT infrastructure with neighboring municipalities. In total, we interviewed 29 practitioners occupying 4 kinds of roles (Table 1) belonging to 16 organizations providing IT services to 54 municipalities (about 17% of the total population of Dutch municipalities). See Appendix A for more information on respondents.

Table 1: **Respondent roles**

Roles	Respondents ($n = 29$)
System Administrator	9
Network Administrator	2
Security Engineer/Officer	12
(C)ISO	6

Pre-interview engagement. Our interviews were one part of a multi-step engagement with the responding municipalities. We first reached out to municipalities via the CERT contact points. We explained the purpose and design of the study and asked them to participate in an interview. The interview had two main purposes: to verify the validity of the detected vulnerable hosts and to understand the reasons for those hosts being present in the network. Discussing and validating specific hosts meant we needed to enable the municipalities to prepare for the interviews. First, we updated our scan results shortly before the engagement. Then we sent a list of hosts, services, and versions that we had detected to our main contact point at each municipality. The contact points then had to identify within their organizations which IT practitioners were responsible for the specific hosts that we wanted to discuss. The practitioners, in turn, would then be able to prepare for the interview, e.g., by checking the exact versions of the services running on the hosts.

Adapting the Research Design. A surprise emerged during the first few interview preparations. Some respondents told us that the hosts we had sent them were not known to them or were not their responsibility, even though they were located in the IP ranges that the municipality had registered with the CERT. They are responsible for updating the data about the corresponding IP ranges if changes happen, yet we consistently ran into discrepancies. We see these discrepancies as outcomes of our research and will discuss this issue in depth in Section 5. That being said, it also meant we had to adapt our research design.

Our first adaptation was to ask respondents to share with us the IP ranges that they were responsible for. This brought another surprise: some of these ranges were completely outside of the ranges that were registered with the CERT. This led to a second adaption: we asked all respondents to share with us, well before the interviews, the IP ranges they were responsible for. We would then scan these ranges, in addition to the ranges that the CERT had on record, for hosts and services in Censys and Shodan. We would also identify vulnerable hosts in these ranges. In this way, we could share up-to-date scan results in preparation for the interview, as well as ensure that we could discuss vulnerable systems. These additional scans brought into focus potentially vulnerable systems that were not attributed by the CERT to the municipality.

Interview protocol. The combination of conducting interviews about actual scanned vulnerable hosts provides an empirical basis for determining how practitioners manage vulnerable systems and reduces the risk of social-desirability bias that may occur with self-reporting. Previous works [13, 33] used qualitative data but did not relate that data to actual systems.

The interviews consisted of three parts. First, we would ask the interviewee to confirm or reject the version information we had inferred about the selected hosts and services. In other words, it acted as a ground-truth validation protocol for the vulnerable services and determined if the external measurements were indeed correct. Second, it seeks to understand practitioner perspectives on the responsibility of IP addresses for monitoring and mitigating vulnerable systems of an organization. Third, it determines how practitioners assess and mitigate vulnerable systems in practice. These discussions centered around the vulnerable systems that we measured, not hypothetical cases. For the latter part, we chose a semi-structured protocol because we wanted the interviewees to freely express their thoughts on the reasons for the presence of vulnerable hosts [25, 29]. The full interview protocol is included in Appendix B.

The interviews were conducted in person or using video conferencing applications and typically took about an hour. There were 6 interviews done with a single respondent, 5 with two, and 4 with three respondents. In those cases, respondents stated that several people were involved in administering the infrastructure and to improve our understanding, all should provide input. One municipality answered interview questions via e-mail because the infrastructure administration required too many people for one interview. In the pre-interview communication, the respondents were informed about the goals of the interview and received an Informed Consent form. At the start of the interview, we reiterated the research goals, gathered consent statements, and asked permission to record the interview for transcription. Respondents participated voluntarily and did not receive compensation for the interview.

Coding. Interviews were transcribed and coded using the ATLAS.ti software [4]. Initial codes were iteratively developed by the lead researcher and two other researchers. The codes clustered topics that described the various kinds of answers of participants. First, 4 interviews were coded by the lead researcher for the initial codes. The research team then refined the initial codes. These codes were then shared with another researcher to independently code a subset of interviews and discuss the results. We refined the codes with the research team with those results, leading to the final codebook. The process of meeting with authors and discussing and independently refining the codebook is a suitable way to ensure the reliability of findings, according to McDonald et al. [36]. The final codebook is available in Appendix C.

5 Validating Observed Vulnerable Systems

The first potential explanation for the presence of vulnerable hosts is that their detection might contain false positives. That is, the network scan data received and disseminated by CERTs might not be fully accurate. Our approach, extracting CPEs from version information in banners, is a normal industry practice, so our data is similar to the data CERTs receive. There is a second type of scanning that does not rely on version information alone and instead uses benign exploit code to test the presence of a vulnerability. Because of the intrusive nature of such scans, we did not adopt this approach. Clearly, the second approach offers greater reliability for estimating vulnerabilities but the approach has a direct impact on the target network, which was unfeasible for our collaborations.

There are two caveats to the approach we used. First, administrators may hide the version in the banners they expose. Second, a service may run a backport, i.e., an older version that includes security patches from a new version but still shows the old banner. A banner rarely shows the presence of a backport of a service. Therefore we asked respondents if they used them. None of our 29 respondents indicated that they (knowingly) ran backports. In two interviews, respondents stated they use a security product that hides version information.

Our scans resulted in 3,402 records from Censys, i.e., that is the number of responding hosts. Within this dataset, 578 hosts were found with at least one versioned service application (17%). A host can run multiple (vulnerable) services on different ports. We, therefore, examine all unique vulnerable CPEs on all hosts. From the 578 hosts, we derived 101 unique CPEs. Of these 101 unique services, 70 contained 1 or more CVEs, with a total of 643 unique CVEs. The 70 vulnerable services were observed in 154 unique hosts in 94 different municipalities. In our population, vulnerable hosts most frequently ran vulnerable versions of OpenSSH and Apache Httpd.

To verify the services and versions at vulnerable hosts, we prepared a list of detected systems for each of the 54 municipalities, based on the IP ranges registered with the CERT. In total, this set contained 24 vulnerable hosts for the 16 IT organizations servicing 54 municipalities.

However, as explained in Section 4.2, the pre-interview communication revealed that the municipalities did not consider some of the CERT-registered IP ranges as falling under their responsibility. So some of the detected vulnerable systems were not under their control or even unknown to our interviewees. This meant that in 9 of the 16 interviews, there were no vulnerable hosts at the municipality in the IP ranges reported by the practitioners. This misalignment is explored further in Section 6. In the remaining 7 interviews, we were able to validate 15 vulnerable hosts. We supplemented this set by validating the 27 non-vulnerable hosts since that inference (vulnerable or not) is based on the exact same data and analysis. This allowed us to test whether using version information from banners is reliable or plagued by a substantial false positive rate. The 15 vulnerable hosts ran 4 different vulnerable services: Apache, OpenSSH, PowerDNS, and Nginx (see Table 2). We confirmed the service and the version with the respondents. We agreed with respondents not to share the versions of the software so as not to facilitate attackers. For all the vulnerable services that we observed with the passive scan data, the actual running service and version were the same, according to the respondents. We also confirmed 27 Microsoft services: 1 MS Internet Information Services 8.0 service, 5 MS Internet Information Services 8.5 services, 9 MS Internet Information Services 10 services, and 12 MS HTTPAPI 2.0 services. Again, all versions were confirmed by the respondents. In total, all 42 observations were correct. While the measurements of the versions of the Microsoft systems are correct, the measurements do not provide insights into their actual vulnerability. First, we cannot determine CPEs that can be linked to specific CVEs. Second, we cannot actively verify vulnerabilities by exploiting them, as described in Section 4.1. During the interviews, we verified the versions but we did not discuss specific security updates. Table 2 provides an overview.

Table 2: **Validation of CPEs identified by Censys**

Service	# Correct	# Incorrect
Apache	6	0
OpenSSH	4	0
PowerDNS	3	0
Nginx	2	0
IIS 8	1	0
IIS 8.5	5	0
IIS 10.0	9	0
HTTPAPI 2.0	12	0
Total	**42**	**0**

Table 3: **The number of hosts, number of versioned hosts and number of vulnerable hosts that the CERT attributes to a municipality, and the practitioners at the municipality themselves.**

Muni Id	Hosts Muni	Hosts CERT	Versioned Hosts Muni	Versioned Hosts CERT	Vulnerable Hosts Muni Only	Vulnerable Hosts CERT Only	Vulnerable Hosts Shared	Total Vulnerable Hosts	Total Unique CVEs
1	26	0	5	0	3	0	0	3	1
2	3	6	0	3	0	2	0	2	45
3	19	13	11	11	0	0	2	2	3
4	22	30	1	1	0	1	0	1	53
5	11	25	3	8	0	0	4	4	14
6	31	47	2	23	0	4	1	5	88
7	21	12	3	5	1	4	0	5	84
8	18	10	3	6	2	5	0	7	73
9	77	156	26	34	0	0	2	2	37
10	24	16	0	0	0	0	0	0	0
11	24	81	0	0	0	0	0	0	0
12	26	50	1	4	0	1	0	1	57
13	13	7	4	0	0	0	0	0	0
14	113	29	8	6	0	1	0	1	31
15	26	0	2	0	0	0	0	0	0
16	46	1	5	0	0	0	0	0	0
Total	500	483	74	101	6	18	9	33	486

6 Attribution and Unclear Responsiblities

As discussed in Section 4.2, our research approach led us to discover a new explanation for the presence of the vulnerable systems: the misalignment between the IP ranges that the municipalities registered with the CERT and the IP ranges that the respondents thought they were responsible for. Perhaps some vulnerable systems persist because the municipalities do not see them as their responsibility, even though they might be notified about their vulnerable status.

In this section, we further explore this misalignment. First, we briefly describe the role of the CERT and the crucial role that the registered IP ranges have in the institutional incident response system. Next, we analyze the relationship between the IPs registered with the CERT and the IPs that the practitioners provided. Then, we quantify the misalignment of the IPs between the CERT and the municipalities. Finally, we examine perspectives on responsibilities.

6.1 Role of CERT and Municipal IPs

The IBD-CERT supports municipalities with security advice and liaises between municipalities and the national CERT. On behalf of municipalities, it contributes to the Baseline Informatiebeveiliging Overheid (BIO) – the Dutch compliance framework for information security within government. One of the goals of the CERT is the detection of incidents and crisis situations and the sharing of knowledge between municipalities and suppliers. While municipalities are ultimately responsible for monitoring their own systems, in pursuing its goals, the CERT also monitors the Internet-facing infrastructure of municipalities. In doing so, it also receives information about vulnerable systems from other parties, such as the national CERT, their own scans, ethical hackers, Shodan, and the Dutch Insititute for Vulnerability Disclosures (DIVD).

6.2 Misaligned Threat Landscape

The misalignment in monitored IP addresses translates to misalignment in the perceived threat landscape by the CERT and the practitioners at the municipality. We find that the CERT generally observes more versioned and vulnerable hosts than the practitioners at the municipality itself. Table 3 describes the number of hosts, number of versioned hosts, and number of vulnerable hosts that the CERT attributes to a municipality, and the practitioners at the municipality themselves. It shows that respondents observed 6 vulnerable hosts that the CERT did not observe. The CERT observes 18 vulnerable hosts that the respondents do not. The respondents and CERT both observe 9 vulnerable hosts. In total, there are 33 vulnerable hosts for the municipalities.

The data in Table 3 shows that neither organization observes the full set of vulnerable hosts. The CERT observes more vulnerable hosts than the practitioners at the municipalities themselves. This would lead the CERT to send notifications to the respective municipalities, who, in turn, would not recognize the system. But the CERT also has a blind spot, the vulnerable systems in the public IP range of the municipality that the CERT does not monitor. In those cases, the CERT could not exercise its supportive function, and would not send any notifications at all.

The differences in vulnerable hosts between the organizations impact the CVEs associated with a municipality. We find that the 15 vulnerable hosts at municipalities result in 62 unique CVEs within IP addresses for which they consider themselves responsible. The vulnerable systems observed by the CERT for those municipalities result in a much larger set: 481 unique CVEs. It is not just the sheer number of vulnerable hosts that is larger. We also see a remarkable difference when we look at the CVSS rating of the vulnerabilities (a score between 1 and 6.9 is considered "low/medium", between 7.0 and 8.9 is "high" and between 9.0 and 10.0 is "critical") [41]. We find that there are 15 critical CVEs in the IP space identified by the respondents versus 107 critical CVEs in the IP space that are registered with the CERT. Similarly, we find that there are 27 high CVEs in the municipality IP space and 191 highs in the CERT IP space. Table 4 depicts the CVEs by severity per organization. The key consequence of the misalignment of the IP ranges is this: the CERT observes much greater risks for the municipalities, compared to the municipalities themselves.

Table 4: **Number of vulnerabilities by CVSS3 severity, as observed by each organization. In parentheses are the number of unique IPs with a vulnerability. Note that most IPs run a service with multiple CVEs of different severity levels, skewing the individual and total IP count.**

Org.	Critical	High	Medium/Low	Total
Muni	15 (7)	27 (14)	20 (9)	62 (15)
CERT	107 (18)	191 (27)	183 (22)	481 (27)

This finding can explain two things. First, it means that the CERT – and in its wake, the national CERT and central government agencies – see the municipalities as much more vulnerable than the municipalities' security practitioners see themselves. That can translate into the perception of the CERT and other government entities that municipalities show a lack of urgency about these issues. Second, it means that the vulnerable hosts persist because the vulnerability notification system is broken. The municipality receives the notifications, but the bulk of these are deemed to fall outside their scope of responsibility. Conversely, there are vulnerable hosts in the self-reported IP ranges that are not registered with the CERT. The municipalities will not be notified about those. Both scenarios lead to vulnerable hosts persisting over time.

During the interviews, most respondents stated that they do not have a complete or up-to-date overview of all the systems that the organization is running externally, so outside the ranges they feel responsible for, but inside the ranges registered with the CERT. Some respondents stated that the internal processes at the municipality for departments to report external systems that they use to the IT or security team were unclear. One respondent (#12) stated "we have a view on those systems that are reported. But not on the ones that are not reported".

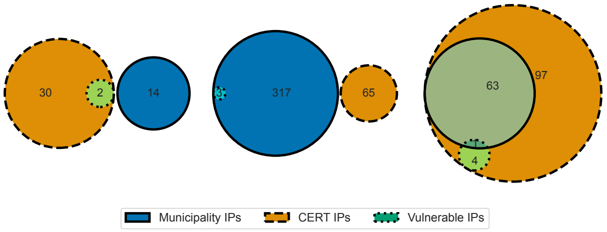

Figure 2: **Three examples of IP sets registered with the CERT and the IPs used by a municipality.**

6.3 Quantifying the Misalignment

To quantify this misalignment, we compare the *sets* of IPs that the CERT associates with a municipality and the IP addresses that the respondents reported as being responsible for. To measure the similarity of those two sets, we use the Jaccard similarity coefficient and the Szymkiewicz-Simpson coefficient, also known as the Overlap coefficient. Both scores range from 0 to 1. The Jaccard score tells us how similar two sets are, where 1 means that all the items appear in both sets. However, the score does not capture when one set is much bigger than the other. Therefore, we also include the Overlap coefficient, which is 1 if one set is fully subsumed in the other set. We find that, on average, the Jaccard coefficient is 0.21 and the Overlap coefficient is 0.56. This means that most municipality IP ranges and CERT IP ranges for that municipality are of a very different size, and also have only partial overlap – meaning, they monitor different addresses. More importantly, the incident response and vulnerability notification infrastructure assume they are 1.

We can also visualize the misalignment using a Venn diagram. Figure 2, shows three examples of sets of IP ranges with vulnerable IPs from the CERT and the municipality itself. The remaining diagrams are in Appendix D. Clearly, the IP ranges registered with the CERT, and the IP ranges used by our respondents are very different.

6.4 Perspectives on Responsibilities

We tried to understand what explains the differences in the attribution of vulnerable hosts. Respondents made a distinction between systems they administer and systems that are used by the organization but are not administered by the respondents, such as SaaS services.

Administering Systems Themselves. Most respondents stated that their responsibility was primarily the systems they administer themselves, and exposing them through their public IP range – i.e., the IPs used for routing Internet traffic in and out of their internal network. They minimize their Internet-facing footprint because it provides the administrators with two clear advantages. First, the "front door" is small and easier to manage for "administrator"-type duties – it gives them a better overview. As one respondent (#2)stated, "the more external points you have, the harder it becomes and you quickly lose oversight". Second, the security tools only need to monitor a small set of addresses which reduces the capacity needed for monitoring those systems. One respondent (#5) stated, "Monitoring simply takes a lot of time, and then we are not even taking into account any type of response". Similar to [3], we found that many municipalities faced staffing issues and limited resources. Consequently, a lot of decision-making prioritized optimized use of (human) resources.

Not Administering Systems. The challenge of overseeing systems is exacerbated when external parties provide a service for the municipality. Most respondents observed a trend in the growing number of Software-as-a-Service (SaaS) platforms. They applaud this trend as it reduces the number of applications and servers they need to manage on-premise. This allows practitioners to handle more work with fewer people. However, the downside is that practitioners are not always aware of those systems, and they no longer administer or control the systems on which their data now resides.

Administering systems is difficult for respondents because of the sheer number of business services a municipality offers, each with its own software application. It is this wide array of services that distinguishes it from a regular enterprise according to one respondent: "A municipality has so many types of connections, partners, and disciplines that it's not comparable to a business. There, they want a single solution but that doesn't exist in municipalities. It's very complex. For example, you could get a notification that the tile of a curb is positioned wrong. The team responsible for that has its own application. We have many demands and wishes and we have a great number of running tenders. Unfortunately."

A majority of respondents stated that there was no clear overview of the platforms that the municipality uses, as those platforms were not always reported to the IT team. For some respondents, a clear process to register newly connected platforms or a mandate to enforce the registration obligation at the IT team was lacking. This is the well-known problem of "shadow IT", which has been plaguing IT managers for decades now [48]. In only two cases did the respondents know the vulnerable host at the IP address that we shared with them, even though they were not responsible for administering it and it wasn't in their public IP range. They knew it was running at a third party. Respondents also brought up the issue of control. They indicated that while procedures are in place to demand security measures during the tender, they do not have the power to exert control. For example, a respondent stated that monitoring – i.e., scanning – external systems was frequently not allowed by the external vendor. Another respondent mentioned that the municipality wanted to ingest logs from the vendor into their SIEM, but the vendor did not want to share its logs, one of the voiced concerns was the privacy of (non-municipality) user information. In all of those cases, the practitioners could not exercise control over how their data was managed at the partner organization, yet they feared that the (political) fallout in case of an incident would be their responsibility.

Administering Shared Resources. Some municipalities collaborate in a governance structure to share ICT resources. This governance structure streamlines resources to reduce costs and optimize the capacity of scarce IT personnel. However, it is not without problems. Respondents in such a governance structure stated that it is hard to draw the line between responsibilities. For example, in two interviews, we spoke to respondents about where the line was drawn between the hardware layer and the application layer. The idea behind this distinction is that the municipalities themselves can manage much of the application layer to accommodate their specific business needs. However, in practice, these layers are often intertwined, and it becomes unclear who is responsible for mitigating a vulnerable system. We observed one case where the organization that runs the infrastructure knows about vulnerable systems in the application layer but is unable to patch the system, mainly because they cannot oversee the impact of an update on the systems that provide the business services.

7 Managing Vulnerable Systems

The last explanation for exposed municipal hosts looks at the reasons that organizations might not patch a system they are responsible for. Practitioners may be unaware of vulnerable systems, or they might be unable to patch them, or they might have decided that patching is not needed, e.g., because they have specific mitigation strategies in place, such as firewall rules or monitoring.

We identified 15 vulnerable systems for which the practitioners we interviewed considered themselves responsible. We asked these respondents how they dealt with those systems. We asked for their rationale on patching, why the systems were not patched, and what if any, other actions had been taken. Due to the nature of our collaboration, we did not learn of the specifics of the business service that a vulnerable system provided. That said, in general, practitioners mentioned that most of their systems contain valuable data, such as personally identifiable information, albeit frequently fragmented. Every bit of valuable data or provided service is considered important, and compromise would have privacy implications even if it applied to only a handful of citizens.

Table 5: **Features of the municipalities, the number of total vulnerable hosts, and the Jaccard and Overlap indexes with CERT. Most municipalities collaborate and provide IT services for several municipalities and auxiliary organizations.**

Muni ID	Vulnerable Hosts	Inhabitants	IT team	Security Team	Servers	Muni's	Aux orgs	Jaccard	Overlap
1	3	10-100k	10-20	yes	100-200	1-2	0	0	0
2	2	10-100k	1-10	yes	0-100	1-2	0	0	0
3	2	100-200k	20-30	no	100-200	5-10	1-2	0.36	0.56
4	1	10-100k	1-10	yes	100-200	2-5	>3	0.038	1
5	4	10-100k	1-10	no	100-200	1-2	0	0.77	1
6	5	200-300k	30-40	yes	>1,000	>10	1-2	0.38	1
7	6	100-200k	20-30	yes	500-1,000	1-2	0	0.16	0.46
8	7	100-200k	10-20	yes	100-200	1-2	0	0	0
9	2	300-400k	30-40	yes	>1,000	5-10	1-2	0.98	0.99
10	0	200-300k	30-40	yes	200-500	1-2	0	0.17	0.79
11	0	100-200k	30-40	no	500-1,000	1-2	0	0.27	1
12	1	10-100k	1-10	yes	100-200	3-5	0	0.47	1
13	0	200-300k	40-50	no	>1,000	3-5	0	0.10	0.18
14	1	100-200k	20-30	yes	200-500	1-2	0	0.02	0.96
15	0	200-300k	40-50	yes	>1,000	5-10	1-2	0	0
16	0	100-200k	40-50	yes	>1,000	1-2	1-2	0	0

7.1 Identifying Vulnerable Systems

We asked respondents how they identified vulnerable systems in their daily work. All respondents stated that they use vulnerability scanners to do so. In addition, they stated that yearly penetration tests are conducted. To stay up to date on the latest vulnerabilities (that might not be incorporated in the vulnerability scanner), they receive security advisories from the CERT, vendors, and popular security news sources. Our findings on the ingress of vulnerability information are in line with earlier work [12], in that practitioners relied on curated vulnerability information from authoritative sources to consider vulnerability information. As we will discuss below, for some vulnerable systems, sysadmins did not know the version of the software they were running. So, they relied on external triggers to become aware of vulnerabilities.

There is a compliance framework in place to act on security advisories. As stated in Section 6.1, the Dutch government uses a compliance framework, BIO, to improve and measure security practices. It contains a chapter on vulnerability management which states that if the severity level of security advisory of the national CERT is marked with probability as "High" and impact as "High", (also known as a "High/High"), the vulnerability should be resolved or mitigated as soon as possible and at the latest within a week. All respondents noted that if they received a "High/High" advisory, they would move to action almost instantaneously.

7.2 Vulnerable vs. No Vulnerable Systems

We analyzed the interviews to determine if municipalities that did not have vulnerable hosts did anything differently than those with vulnerable hosts. We included 5 municipal IT organizations where we detected no vulnerable hosts. Remember that our measurement approach relies on obtaining versions from banner data. If administrators hide the version information, our method will not determine vulnerable hosts. We learned that 2 of the 5 municipal organizations indeed ran security products that obscured versions.

At 11 municipal organizations, we did observe vulnerable hosts, either in the IPs registered at the CERT or administered by the municipality. We tried to find a common denominator for organizations with vulnerable hosts versus those without. We looked at security tools, the capacity of IT staff, the size of the municipality, the size of the IT team, the presence of a security team, and the act of vulnerability scanning.

All the respondents we interviewed, with and without vulnerable hosts, had basic generic security tools in place, like firewalls, NAT, EDR, logging, and network segmentation. Similarly, all respondents noted that they did not have sufficient capacity, in terms of qualified IT staff, for their security duties. Next, we checked if the size of the municipality, measured by the number of inhabitants, had a relation to the number of vulnerable hosts. We observed that both the smallest and largest municipalities had vulnerable hosts, as did several in between. We then compared the size of the IT team but also found that vulnerable hosts occurred in small and bigger IT teams. Finally, we looked at organizations in terms of the number of servers they ran. These servers are not all exposed to the Internet, but the number of app servers acts as a proxy for the size of their infrastructure. The idea is that a larger infrastructure might contain more vulnerable machines. However, here too we see no clear distinction. Next, we observed that all respondents, except one, engaged in vulnerability scanning, discounting this as an explanation for the presence or absence of vulnerable hosts. The one respondent who did not actively use it was not part of the security team. Finally, we wondered if the presence of a security team might have an effect. From the 16 interviews, 12 organizations ran a security team, and 4 did not. We observed that 9 municipalities with a security team have vulnerable hosts. Of the 4 municipalities without a security team, 2 municipalities exposed vulnerable hosts.

In sum, we do not observe a link between specific features of an organization and the number of vulnerable hosts. We also do not find substantially different security practices among respondents in the interview data.

7.3 No Patch

We hypothesized that observed vulnerable systems are indeed vulnerable, but administrators may have their reasons for not patching. We first examined if they had put specific mitigation strategies in place for the vulnerable hosts. All the respondents stated that they had generic mitigation strategies in place (such as network segmentation, logging, firewalls, and Intrusion Detection Systems). Some respondents stated that they also use a managed SIEM. However, only in one case was a specific mitigation strategy in place for an observed vulnerable system. That system was run on an isolated network. Without specific mitigation strategies in place, we asked administrators what other actions, if any, were taken for the vulnerable hosts. We analyzed the interviews for the 15 vulnerable hosts and synthesized the responses into various explanations that we condensed into four categories: *unaware* of the vulnerable system, *unable* to patch it, the system was (in the process of being) *retired*, or the system was *shut down*. We describe each in more detail below, and tabulate an overview of these explanations in Table 6.

Table 6: **Explanations for vulnerable systems**

Explanation	**# Systems**
Unaware	9
Unable	3
Retired	3
Shut down	N/A
Total	**15**

Unaware. We observed during the interviews that respondents were not always aware of the vulnerable system. We encountered three types of unawareness. First, an administrator retired the system and assumed it was no longer online. This, however, was not the case. One respondent (#6) said: “this is a system that is phased out. I’m actually surprised about this. Thanks, I’ve got some homework to do.” Second, there was a case where a vulnerable system was not on the radar at the organization. When presented with the system, the security team could not find the system in their asset inventory but acknowledged it was running in their IP range. There was no direct explanation for that situation. Third, the system was not directly identified as a vulnerable system. The respondent (#1) stated: “These servers run directly from the Debian repos. We did not patch these systems because we assume the repo provides a decent package”. For this particular case, the administrator was not part of the security team, so it may have been flagged as vulnerable elsewhere in the organization. However, the actual administrator of the system initially did not consider it vulnerable.

We investigated this type of unawareness further by checking if respondents knew about the versions of the software they run and if they were vulnerable. None of the respondents during the interview explicitly had the versions of the software they were running at the top of their minds. Similarly, none of the 29 respondents directly knew the latest version of the software they were running. All had to look up the service and version from their asset inventory system, most often a vulnerability scanning report.

We tried to gauge if respondents knew ex-ante if they set up a vulnerable system in their infrastructure. Most respondents stated that the software that is run is installed to the latest version when set up, regardless of potential vulnerabilities, because that is the best they can do. Potential vulnerabilities will be found when the vulnerability scanner is run. One respondent indicated that when installing the latest version, he checked the latest packages for information and vulnerabilities from the distributor’s repository to make sure a newer version would not be released in the very near future.

Unable. In two cases, practitioners reported that they were unable to patch a vulnerable system. This happened either because of a lack of mandate from the organization or because the vulnerable software was a dependency in a product that was used to provide a business service.

At one organization, the security team was aware of the vulnerable system but was unable to undertake mitigating actions. This particular situation derived from a governance structure where the respondents were part of an organization that was responsible for the majority ICT infrastructure of the municipalities but not the last application layer. The application layer was managed by a small IT team at the municipality itself. The security team at the organization managing the infrastructure is somewhat involved in the management of the application layer due to their expertise, but they did not have the mandate to intervene themselves. Another consideration was that a patch could break the services offered by the vulnerable system. A respondent (#13) stated “fixing vulnerabilities in the application layer is outside the scope of our mandate. That said, it isn’t entirely that black and white. But in this case, we cannot functionally oversee the consequences for the underlying application and business processes.”

At another organization, the vulnerable host ran a product with a dependency that was vulnerable. The vulnerable software could only be patched by updates from the product.

The risks of that system were mitigated by the monitoring of an endpoint security product – that ran on all managed devices – and by running the system on an isolated segmented network. For that host, the security team did not directly intervene in mitigating the risk but had to engage other teams to act on the vulnerable system. The respondent (#12) stated: “that’s why I’m pushing these people to act on this system”. But sometimes such a system could not be removed. The respondent (#12) said: “this particular system, it is provided by a supplier and someone in our organization opted for that product. How do you deal with it? Retiring the system is the only way”. The system was not actually retired because of business reasons.

Retired. In various cases, the observed vulnerable hosts were in the process of being retired. A retired system should not be online, but there is some time between deciding to retire the machine and it actually being offline.

Retiring a system could have various reasons but, most often, respondents stated that it was a (legacy) system that is outdated. For example, some systems ran outdated software, and the service it provided can now be done better with a new product. Consequently, the system was phased out and didn’t get much attention anymore. The legacy system was marked for retirement and lingered around for a while before it actually went offline. A respondent (#6) said “this was a great product at the time but these days it’s outdated. We are now in the phase of retiring this system”. In another case, we found a vulnerable system before the interview, and during the interview, the system was no longer online, the respondent (#23) stated that: “by now that system is turned off”. This happened only once.

Shut down. One of the most drastic mitigation strategies for vulnerable systems is simply “pulling the plug”. This was not done for any of the vulnerable systems we observed, but many respondents stated that shutting a system down was part of their toolbox of mitigation strategies – albeit one very few actually want to use. When asked if this was actually done in practice, only a few respondents stated that this was within their power to actually execute. The majority of respondents stated that, while it is an option, in practice, shutting down is hardly done because business continuity takes priority. For most respondents, the only situation where they actually shut down systems was during the Log4J vulnerability. The severity of the vulnerability and the fact that it was not clear which software was vulnerable allowed for enough organizational pressure to trump business continuity. As one respondent stated: “once a vulnerability hits the news, people start taking it seriously. Sometimes we need an incident like that to make strides in security”.

But without that sense of urgency, business overpowers security. One respondent (#22) stated, “business and security sometimes have opposing interests. Contrary to a commercial organization, a municipality has certain societal obligations, therefore we simply can’t shut a system down like that because we are required to offer those services. Considerations are complex and discussions are quickly taken out of context. Then it doesn’t matter what actually happened, but how it is perceived because something is in the newspaper”.

7.4 Patching Systems: Prioritizing

Respondents stated that patching is the preferred strategy to deal with a vulnerable system. Sometimes, a patch is not directly available, as was the case with a Citrix vulnerability in 2020 [1], then they rely on the mitigation strategies provided by the vendor or the security advisory. However, the capacity, in terms of people, for rolling out patches is limited. The vast majority of respondents stated capacity as an obstacle to security. As practitioners face vulnerabilities in their Internet-facing systems as well as their internal networks, this forces them to prioritize what systems to patch first. In doing so, there are two main criteria: a) whether the system is Internet-facing and b) the severity level of the vulnerability.

First, several respondents stated that all vulnerabilities should be dealt with. This self-reporting, however, is contradicted by the fact that we did find vulnerable hosts, illustrating the limitations of self-reporting on patching behavior. To be fair, many respondents also appeared to accept that there will always be vulnerabilities somewhere in their infrastructure. Vulnerabilities in Internet-facing systems should be dealt with first – as they consider it the front door to their internal network. One respondent (#26) stated: “Internet-facing systems have priority because the chance of abuse is higher than systems inside our network.” If this is true, then there should be more vulnerabilities in internal systems versus the Internet-facing systems. We could not measure this directly, but when asked, several respondents admitted that there were indeed (many) more vulnerabilities in their internal network.

While many respondents stated that their Internet-facing systems should not contain vulnerabilities, it was also mentioned that they do not consider it likely that a real attack would happen there. Instead, they feared an attack via an unsuspecting user clicking a link in a phishing mail.

Second, respondents indicated that systems with a high-severity vulnerability are prioritized. The severity metric is most often determined by either the severity level of the security advisory or the severity score of the vulnerability scanner. The Common Vulnerability Scoring System (CVSS) cite [20] is a popular scoring system to determine the severity of a vulnerability. The CVSS score is popular but has its limitations. For example, it does not take into account the ease of exploitability of a vulnerability, the availability of an exploit, or details on the number of exploitations of the vulnerability in the wild. Many security companies expand on CVSS with their own data and ranking to improve the assessment of the severity of a vulnerability. In doing so, (proprietary) vulnerability scanners often report the CVSS as well as their own scoring system. For example, the vulnerability scanner Nessus – popular among respondents – provides its own Vulnerability Priority Rating (VPR). The VPR takes additional factors into account, such as the CVSSV3 Impact score, the age of the vulnerability, the exploit code maturity, and more [50]. If vulnerabilities have a High or above classification (in any kind of scoring methodology), they are quickly prioritized according to the respondents.

To verify this claim for high-severity security advisories, we examined the national CERT security advisories for "High Impact/High Probability" vulnerabilities and referenced them with the systems we scanned. The security advisories contain vulnerabilities for open-source software but also proprietary software. We could not validate the claim that respondents patched "High/High" advisories for proprietary software quickly because the associated software services could not be fingerprinted by us for a version. However, in one case, security company Fox-IT wrote in their blog that they could fingerprint Citrix software for two specific CVEs [22]. In the CERT dataset with IP addresses, we reproduced their method and observed that the systems with CVEs were applicable for 7 hosts. Looking historically, in the two weeks after the advisory was sent and the update became available, these systems were patched, giving credibility to the respondents' claims.

8 Discussion

Explanations for persisting vulnerable systems. We consider three explanations for the 154 vulnerable systems we observed in the IP ranges registered with the CERT: incorrect measurements, misalignment in the responsibility of IPs, and vulnerable machines unpatched for some other reason. We interviewed 16 municipal IT organizations that had 33 of the vulnerable systems.

We learn that the first explanation doesn't explain any of the vulnerable systems. We find that the external measurements of hosts using banner information are not plagued by false positives. The second explanation, the misalignment of IPs – i.e., IPs registered at the CERT and those used by practitioners – is observed at all the 16 municipal IT organizations we spoke, and this most likely happens at many more, if not all, municipalities. Of the 33 vulnerable machines for the municipalities, 18 vulnerable hosts were seen only by the CERT, 6 by the municipalities alone, and 9 were seen by both organizations. This brings us to the third explanation: vulnerable machines that are unpatched for other reasons. Of the 15 vulnerable hosts observed by the municipalities, 9 are explained by administrators being unaware of those systems.

In short, the main explanation is the misalignment in the attribution of IP ranges, where administrators do not consider the systems their responsibility. This explanation is followed at some distance by the explanation that the IT organization was unaware of the presence of the vulnerable systems.

What is causing the misalignment problem? As Vermeer et al. [54] noted, organizations consistently struggle to keep a complete inventory of their assets. The assets are constantly changing, with many changes unplanned or unrecorded. This is closely related to the problem of "shadow IT" – systems and services that are "not known, accepted and supported" by an organization's official IT department [48]. Indeed, when respondents were speculating what the vulnerable hosts were, they frequently mentioned SaaS solutions and specific services contracted by some department of the municipality, but outside their purview, the purview of the IT department. This is classic "shadow IT". This explains why they did not consider it their responsibility to safeguard these systems. It also suggests that this is most likely not a problem exclusively to municipalities, and we might expect this also to occur in other enterprise environments. Not only does it mean there is no clear responsibility to keep those systems secure, but it also gravely undermines the CERT-based notification mechanism. Those notifications reach the IT department, but cannot find their way to the actual entity managing the host, because IT does not know. The fact that the IT teams see themselves as powerless towards "shadow IT" does not reduce the actual risk for the organizations. The vulnerable systems continue to run, exposing 183 "Low/Medium", 191 "High", and 107 "Critical" CVEs, according to IPs from the CERT data.

Finally, our findings also highlight a discrepancy between the widely-held view that local governments are very vulnerable, as mentioned in the Introduction, and the perspective of the local IT departments, as the latter observe far fewer problems in their own systems. This discrepancy might explain why the current situation persists, even though CERTs and higher levels of government keep warning local governments.

The municipal IT departments do their work under serious resource constraints. The lack of capacity and staff frequently came up. We observed that vulnerable hosts, with the exception of one, did not have specific mitigation strategies in place. Instead, respondents rely on generic mitigation measures. Respondents frequently mention the lack of capacity for IT security tasks. This aligns with our finding that organizations have only generic security measures in place – a rational choice to maximize defenses with limited resources. Coincidently, in December 2022, the Association of Dutch Municipalities (VNG) requested additional resources from the central government to increase its IT capacity in support of the Dutch National Cybersecurity Strategy, but those resources were denied [30]. Simply put, cybersecurity requirements increase, yet resources to comply lag behind.

Recommendations. So, what can municipalities do to tackle the attribution issue? After all, they themselves registered the IP addresses with the CERT. Clearly, they would benefit from keeping the IP ranges registered with the CERT up-to-date. At a minimum, this allows notifications to reach the correct entity. If they are unaware of vulnerable systems, the notifications should help address that. What appears needed is some guidance or support on how to handle the responsibility gap that exists around “shadow IT”. Who is responsible for what system? The owner of the system? The owner of the data in the system? Respondents occasionally stated that they feel somewhat responsible for systems running elsewhere since they handle municipal data. At the same time, another respondent stated that they are not allowed to scan the systems of their partners, so they are unaware that their data resides in a vulnerable system. Until a clear delineation of responsibilities exists, organizations remain at risk for cyber threats. The new NIST 2.0 framework may guide practitioners in this respect. It emphasizes a new Govern function to gain a better “understanding of cybersecurity roles and responsibilities” [39].

For CERTs, our findings suggest two points. First, the institutional framework to mitigate cyber threats assumes that vulnerable systems are correctly attributed. This assumption turns out to be problematic. The national CERT detects a vulnerable system and notifies the relevant sectoral CERT. In turn, the sectoral CERT notifies the appropriate constituent about the vulnerable system. The constituent is then expected to take action. But in practice, we observe that the last, crucial step of this notification process is flawed – as the owner of the vulnerable system is often not the recipient of the notification. As a result, we find that those systems remain vulnerable, with the data of the organization at risk. This finding contests the effectiveness of the institutional framework for notifying vulnerable entities and highlights the need for a better reporting process between the CERT and its constituents.

Second, for CERTs, like for municipalities, there seems to be a need for a clear delineation of responsibilities for external systems. Currently, what constitutes the network of an organization appears to be diffuse, but counting only on-premise infrastructure seems archaic. Perhaps a recommendation could be to build a feedback loop for the notifications, a bit like a ticketing system. This way, if a notification is not picked up by the entity to whom it is assigned, because that entity sees it as outside its responsibility, then the ticket gets returned to the CERT. Both the CERT and the municipal IT leadership, e.g., the CISO, can then observe what systems are vulnerable, yet not acted upon. This is then a starting point for identifying who is managing those systems. Currently, neither the CERT nor the municipalities seems aware of the scale problem we have uncovered. If notifications function like tickets, then the scale and location of the problem become very clear to see for all parties involved.

Limitations. Our research design introduces several limitations: external validity, internal validity, scope, and desirability bias. First, we focused exclusively on Dutch municipalities. This risks that our findings might not be generalizable. We believe, though, that CERTs worldwide face similar challenges in delineating responsibilities and correctly attributing vulnerable hosts to their constituents for effective notifications and subsequent remediation. Our findings point to the problem of “shadow IT”: the challenges of managing asset inventory, including IPs and externally run services, which is a problem that many other organizations deal with. Also, a sample size of 16 interviews is limited. This sample, however, does cover 17% of the total population of municipalities. Second, our respondents were directly involved with managing vulnerable systems. Yet, they are only part of the (security) IT teams within the organizations, and their knowledge might be incomplete. We allowed additional respondents during the interviews to mitigate this risk. Yet, the limited knowledge is not just a barrier to measurement, but also an operational reality with consequences: respondents did not know who was responsible for large portions of the IP ranges that were registered with the CERT. Hence, they cannot delegate vulnerability notifications, let alone ensure mitigation.

Third, we could only discuss Internet-facing systems that ran a service with a version. Most proprietary (security) products run a service that does not include (backported) version information in the banner. Thus, the number of exposed vulnerable hosts is likely to be higher. That said, our methodology is limited as we could not perform active measurements of observed vulnerabilities. Therefore, the measured vulnerability of systems does not directly map to an equal amount of risk.

Lastly, the interviews risk desirability bias. When asked why systems were vulnerable, participants might give answers to paint them in a favorable light. We tried to mitigate this risk by a) talking about actual systems, b) interviewing respondents without their superiors, c) keeping the interviews confidential for the municipality and the CERT, and d) stating in the interview that we were scientists and were not there to pass judgment. We also believe that our methodology, interviewing respondents about actual systems, led us to a new discovery of the misalignment in responsibilities for vulnerable systems. Without our approach, practitioners would have reported very few, if any, vulnerable systems in their networks.

9 Conclusion

We asked what explains unpatched vulnerable systems that are detected and notified about, but not patched. We found that the detections are correct. It turned out that most of the vulnerable systems fell into a responsibility gap around “shadow IT” and other systems outside the reach of the IT organizations. Sysadmins did not consider these systems their responsibility. The CERT was not aware. We further identified that for most systems that did fall under the sysadmin’s responsibility, they were unaware of the presence of the vulnerabilities. Our findings highlight the need to re-evaluate and improve the critical institutional structures of incident response and vulnerability notifications.

Acknowledgments

We are grateful for our collaboration with IDB-CERT, which allowed us to connect with Dutch municipal IT-teams. We thank the anonymous reviewers for their constructive feedback that improved the paper. This work was supported by the Ministry of the Interior and Kingdom Relations of the Netherlands and Delft University of Technology under Grant M75B07 and partially by NWO-project "THESEUS" (NWA.1215.18.006).

References

[1] Lawrence Abrams. Citrix ADC CVE-2019-19781 Exploits Released, Fix Now! Online: https://www.bleepingcomputer.com/news/security/citrix-adc-cve-2019-19781-exploits-released-fix-now/, November 2020. Accessed: 2023-06-01.

[2] AfricaCERT. Mission statement. Online: https://www.africacert.org/mission-statement/, 9 2023. Accessed: 2023-09-6.

[3] Noura Alomar, Primal Wijesekera, Edward Qiu, and Serge Egelman. "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)*, pages 319–339, 2020.

[4] Atlas.Ti. ATLAS.ti | The #1 Software for Qualitative Data Analysis. Online: https://atlasti.com. Accessed: 2023-03-01.

[5] Rob Barrett, Eser Kandogan, Paul P. Maglio, Eben M. Haber, Leila A. Takayama, and Madhu Prabaker. Field studies of computer system administrators: analysis of system management tools and practices. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, CSCW '04, pages 388–395, New York, NY, USA, November 2004. Association for Computing Machinery.

[6] T. G. Bondar, Hala Assal, and Abdelrahman Abdou. Why do Internet Devices Remain Vulnerable? A Survey with System Administrators. *Proceedings 2023 Workshop on Measurements, Attacks, and Defenses for the Web*, 2023.

[7] Censys. Frequently Asked Questions. Online: https://support.censys.io/hc/en-us/articles/360038378552-Frequently-Asked-Questions, January 2023. Accessed: 2023-03-2.

[8] CERT.br. About CERT.br. Online: https://www.cert.br/about/, 9 2023. Accessed: 2023-09-6.

[9] F. Cetin, C. Gañán, Maciej Korczyński, and M. V. Eeten. Make notifications great again: learning how to notify in the age of large-scale vulnerability scanning. In *WEIS*, pages p. 23, 2017, 2017.

[10] Niraj Chokshi. Hackers Are Holding Baltimore Hostage: How They Struck and What's Next. *The New York Times*, May 2019. Accessed: 2023-05-15.

[11] Chris Teale (GCN). Unpatched, known vulnerabilities still key driver of cyberattacks. Online: https://gcn.com/cybersecurity/2023/03/unpatched-known-vulnerabilities-still-key-driver-cyberattacks/383489/, March 2023. Accessed: 2023-05-15.

[12] Stephanie de Smale, Rik van Dijk, Xander Bouwman, Jeroen van der Ham, and Michel van Eeten. No One Drinks From the Firehose: How Organizations Filter and Prioritize Vulnerability Information. In *IEEE Symposium on Security and Privacy (SP) (SP)*, pages 203–219, Los Alamitos, CA, USA, May 2023. IEEE Computer Society.

[13] Constanze Dietrich, Katharina Krombholz, Kevin Borgolte, and Tobias Fiebig. Investigating System Operators' Perspective on Security Misconfigurations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS '18, pages 1272–1289, New York, NY, USA, October 2018. Association for Computing Machinery.

[14] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. A Search Engine Backed by Internet-Wide Scanning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, CCS '15, pages 542–553, New York, NY, USA, October 2015. Association for Computing Machinery.

[15] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, and J. Alex Halderman. The Matter of Heartbleed. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, IMC '14, pages 475–488, New York, NY, USA, November 2014. Association for Computing Machinery.

[16] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. {ZMap}: Fast Internet-wide Scanning and Its Security Applications. In *22nd USENIX Security Symposium (USENIX Security 13)*, pages 605–620. USENIX Association, August 2013.

[17] Dutch Safety Board. Vulnerable through software - Lessons resulting from security breaches relating to Citrix software. Online: http://www.onderzoeksraad.

nl/en/page/17171/vulnerable-through-software---lessons-resulting-from-security, December 2019. Accessed: 2023-05-15.

[18] Harun Ecik. Comparison of Active Vulnerability Scanning vs. Passive Vulnerability Detection. *2021 International Conference on Information Security and Cryptology (ISCTURKEY)*, pages 87–92, 2021.

[19] FIRST. CSIRT Services Framework Version 2.1. Online: https://www.first.org/standards/frameworks/csirts/csirt_services_framework_v2.1, November 2019. Accessed: 2023-05-15.

[20] FIRST. CVSS v3.1 Specification Document. Online: https://www.first.org/cvss/specification-document, March 2019. Accessed: 2023-05-15.

[21] Federal Office for Information Security. CERT-Bund. Online: https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Cyber-Sicherheitslage/Reaktion/CERT-Bund/cert-bund.html?nn=907524, 9 2023. Accessed: 2023-09-6.

[22] Fox-SRT. CVE-2022-27510, CVE-2022-27518 – Measuring Citrix ADC & Gateway version adoption on the Internet. Online: https://blog.fox-it.com/2022/12/28/cve-2022-27510-cve-2022-27518-measuring-citrix-adc-gateway-version-adoption-on-the-internet/, December 2022. Accessed: 2023-05-04.

[23] Béla Genge and Călin Enăchescu. ShoVAT: Shodan-based vulnerability assessment tool for Internet-facing services. *Security and Communication Networks*, 9(15):2696–2714, 2016. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/sec.1262.

[24] Jonathan Greig. State of emergency declared as City of Oakland grapples with ransomware attack. Online: https://therecord.media/oakland-ransomware-emergency-declared, May 2023. Accessed: 2023-05-11.

[25] Dean Hammer and Aaron Wildavsky. The Open-Ended, Semistructured Interview: An (Almost) Operational Guide. In *Craftways*. Routledge, 2 edition, 1993. Num Pages: 45.

[26] IB&P. Kwetsbaar door software - Lessen naar aanleiding van beveiligingslekken door software van Citrix - Informatiebeveiliging & Privacy. Online: https://ib-p.nl/download/kwetsbaar-door-software-lessen-naar-aanleiding-van-beveiligingslekken-door-software-van-citrix/, February 2022. Accessed: 2023-05-24.

[27] Platon Kotzias, Leyla Bilge, Pierre-Antoine Vervier, and Juan Caballero. Mind Your Own Business: A Longitudinal Study of Threats and Vulnerabilities in Enterprises. In *Proceedings 2019 Network and Distributed System Security Symposium*, San Diego, CA, 2019. Internet Society.

[28] Katharina Krombholz, Wilfried Mayer, Martin Schmiedecker, and Edgar Weippl. "I have no idea what i'm doing": on the usability of deploying HTTPS. In *Proceedings of the 26th USENIX Conference on Security Symposium*, SEC'17, pages 1339–1356, USA, August 2017. USENIX Association.

[29] Beth L. Leech. Asking Questions: Techniques for Semistructured Interviews. *PS: Political Science & Politics*, 35(4):665–668, December 2002. Publisher: Cambridge University Press.

[30] Alexander Leeuw. Geen extra geld voor uitvoering cybersecurity, December 2022. Section: digitaal.

[31] Rober Lemos. CISA Addresses 'Cyber Poor' Small Biz, Local Government. Online: https://www.darkreading.com/threat-intelligence/cisa-addresses-cyber-poor-small-biz-local-government, May 2023. Accessed: 2023-05-16.

[32] Frank Li, Zakir Durumeric, Jakub Czyz, Mohammad Karami, Michael Bailey, Damon McCoy, Stefan Savage, and Vern Paxson. You've Got Vulnerability: Exploring Effective Vulnerability Notifications. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 1033–1050, Austin, TX, August 2016. USENIX Association.

[33] Frank Li, Lisa Rogers, Arunesh Mathur, Nathan Malkin, and Marshini Chetty. Keepers of the Machines: Examining How System Administrators Manage Software Updates For Multiple Machines. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 273–288, Santa Clara, CA, August 2019. USENIX Association.

[34] Gordon Lyon. Nmap: the Network Mapper - Free Security Scanner, 5 2023.

[35] John Matherly. *Complete Guide to Shodan*. Leanpub, July 2015.

[36] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–23, November 2019.

[37] Josh Meyer. Local governments are more vulnerable to cyberattacks than ever before. DHS wants mayors to step up. Online: https://www.usatoday.com/story/news/politics/2022/02/08/local-government-cybersecurity-digital-threats/9208951002/, August 2022. Accessed: 2023-05-24.

[38] NCSC.gov.uk. Early Warning - NCSC. Online: https://www.earlywarning.service.ncsc.gov.uk/, 9 2023. Accessed: 2023-09-6.

[39] NIST. Cybersecurity Framework. Online: https://www.nist.gov/cyberframework, November 2013. Accessed: 2023-06-2.

[40] NIST. National Vulnerability Database - Home. Online: https://nvd.nist.gov/, May 2023. Accessed: 2023-05-17.

[41] NIST. National Vulnerability Database - Vulnerability Metrics. Online: https://nvd.nist.gov/vuln-metrics/cvss, May 2023. Accessed: 2023-05-17.

[42] Donald Norris. A Look at Local Government Cybersecurity in 2020. Online: https://icma.org/articles/pm-magazine/look-local-government-cybersecurity-2020, July 2021. Accessed: 2023-05-24.

[43] NOS Nieuws. Hack bij gemeente Hof van Twente veroorzaakt door te simpel wachtwoord. Online: https://nos.nl/artikel/2372868-hack-bij-gemeente-hof-van-twente-veroorzaakt-door-te-simpel-wachtwoord, March 2021. Accessed: 2023-05-11.

[44] Jamie O'Hare, Rich Macfarlane, and Owen Lo. Identifying Vulnerabilities Using Internet-Wide Scanning Data. In *2019 IEEE 12th International Conference on Global Security, Safety and Sustainability (ICGS3)*, pages 1–10, January 2019.

[45] The Canadian Press ·. Ontario police warn of recent cyberattacks targeting local governments | CBC News. Online: https://www.cbc.ca/news/canada/toronto/cyberattacks-targetting-local-government-ontario-1.4824772, September 2018. Accessed: 2023-05-11.

[46] Qualys. Qualys VMDR - Vulnerability Management Tool | Qualys, 9 2023. Accessed: 2023-09-20.

[47] Rapid7. Metasploit, 9 2023. Accessed: 2023-09-20.

[48] Christopher Rentrop and S. Zimmermann. Shadow IT - Management and Control of Unofficial IT. In *Proceedings of the 6th International Conference on Digital Society*, January 2012.

[49] Chris Teale. Southern states have the most open cyber exposures, report finds. Online: https://gcn.com/cybersecurity/2023/02/southern-states-have-most-open-cyber-exposures-report-finds/383418/, February 2023. Accessed: 2023-05-24.

[50] Tenable. CVSS Scores vs. VPR (Nessus 10.5). Online: https://docs.tenable.com/nessus/Content/RiskMetrics.htm. Accessed: 2023-05-2.

[51] Tenable. Nessus Vulnerability Scanner, 9 2023. Accessed: 2023-09-20.

[52] Christine Utz, Matthias Michels, Martin Degeling, Ninja Marnau, and Ben Stock. Comparing Large-Scale Privacy and Security Notifications. In *Proceedings on Privacy Enhancing Technologies*, Lausanne, Switzerland, July 2023. ISSN: 2299-0984.

[53] Nicole F. Velasquez and Suzanne P. Weisband. System administrators as broker technicians. In *Proceedings of the Symposium on Computer Human Interaction for the Management of Information Technology*, CHiMiT '09, pages 1–8, New York, NY, USA, November 2009. Association for Computing Machinery.

[54] Mathew Vermeer, Jonathan West, Alejandro Cuevas, Shuonan Niu, Nicolas Christin, Michel Van Eeten, Tobias Fiebig, Carlos Ganan, and Tyler Moore. SoK: A Framework for Asset Discovery: Systematizing Advances in Network Measurements for Protecting Organizations. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 440–456, Vienna, Austria, September 2021. IEEE.

[55] Daniel Votipka, Rock Stevens, Elissa Redmiles, Jeremy Hu, and Michelle Mazurek. Hackers vs. Testers: A Comparison of Software Vulnerability Discovery Processes. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 374–391, May 2018. ISSN: 2375-1207.

[56] Jonathan Codi West and Tyler Moore. Longitudinal Study of Internet-Facing OpenSSH Update Patterns. In Oliver Hohlfeld, Giovane Moura, and Cristel Pelsser, editors, *Passive and Active Measurement*, Lecture Notes in Computer Science, pages 675–689, Cham, 2022. Springer International Publishing.

A Respondent Details

Table 7: **Overview of the respondents per municipality**

Respondent Id	Muni Id	Role	Gender	Vulnerable Hosts
1	1	System Administrator	Male	3
2	2	System Administrator	Male	2
3	3	System Administrator	Male	1
4	3	Security Officer/Engineer	Male	1
5	4	(C)ISO	Male	1
6	5	System Administrator	Male	4
7	6	System Administrator	Male	5
8	6	(C)ISO	Male	5
9	7	Network Administrator	Male	5
10	7	(C)ISO	Male	5
11	7	Security Officer/Engineer	Male	5
12	8	Security Officer/Engineer	Male	7
13	9	Security Officer/Engineer	Male	2
14	10	Security Officer/Engineer	Male	0
15	10	Security Officer/Engineer	Male	0
16	11	Security Officer/Engineer	Male	0
17	11	System Administrator	Male	0
18	12	System Administrator	Male	1
19	12	System Administrator	Male	1
20	12	System Administrator	Male	1
21	13	Security Officer/Engineer	Male	0
22	13	(C)ISO	Female	0
23	13	(C)ISO	Male	0
24	14	Network Administrator	Male	1
25	14	Security Officer/Engineer	Male	1
26	15	(C)ISO	Male	0
27	16	Security Officer/Engineer	Female	0
28	16	Security Officer/Engineer	Male	0
29	16	Security Officer/Engineer	Male	0

B Interview Protocol

Introduction and background

- Can you tell me about yourself?
- What does a typical day look like?
- How many devices and servers are you managing, and how big is the team?
- How is security organized in your organization?
- What do you consider the biggest obstacles in security?
- What are the Internet-facing systems of the municipality?
- How do you monitor those systems?
- How do you stay up to date on vulnerabilities? Is that an active process?

Advisories and CERT

- Who receives advisories and notifications from the CERT?
- Who manages the IP ranges, an individual or a team?
- Who is responsible for following up after a notification from the CERT?
- How do you determine if a notification is relevant?
- Do you report changes in your infrastructure to the CERT?

Specific Vulnerable Systems

- For system X, we detected service Y and version Z. Is that correct? Did you run a backport? Is it vulnerable? If so, which CVEs? How did you obtain that information?
- Are you aware of the latest version of service X? How do you obtain that information?
- How do you deal with those CVEs?
- Did you apply mitigation strategies? Why?
- Do you have other mitigation strategies that were not used? Do you have examples?
- How does the location or function of the system influence the choice of a mitigation strategy?

C Codes

Determining Vulnerable Systems

Subcodes Active Search; Asset Inventory; Notifications; Not determined; Vulnerability Scanning.

Mitigation Strategies

Subcodes Strategies Notifier; Security Tools; Isolating systems; Managed Services; Non-internet facing.

Prioritization and Patch Practices

Subcodes Critical versus non-critical score; Internet-facing vs non-Internet-facing; Latest version on install.

Responsibilities

Subcodes Compliance; Dependencies in products; External, Cloud and Saas services; Public IPs and internal network; Security role in organization; Users and awareness.

Underming Security

Subcodes Capacity for security tasks; Budget; People and skills; Priorities, partners and collaborations; Legacy systems.

D Venn Diagrams Total IP Sets

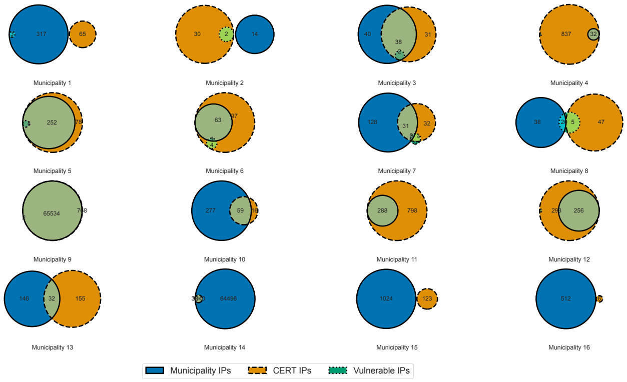

Figure 3: **IP sets registered with the CERT and the IPs used by a municipality.**