Are You Sure You Want To Do Coordinated Vulnerability Disclosure?

Ting-Han Chen
University of Twente
The Netherlands
t.h.chen@utwente.nl

Carlotta Tagliaro
TU Wien
Austria
carlotta@seclab.wien

Martina Lindorfer
TU Wien
Austria
martina@seclab.wien

Kevin Borgolte
Ruhr University Bochum
Germany
kevin.borgolte@rub.de

Jeroen van der Ham-de Vos
University of Twente
The Netherlands
j.vanderham@utwente.nl

***Abstract*—The rising numbers of vulnerabilities and security issues stemming from the rapid iteration and development of the Internet of Things (IoT) have introduced new challenges for the involved stakeholders to mitigate them in time. To effectively bring researchers, vendors, and end-users together to address such problems, Coordinated Vulnerability Disclosure (CVD) has become standard practice. Although general CVD procedures for practitioners to follow exist, adapting them to the specific circumstances has proven to be complicated in practice.**

In this paper, we document our experience of reporting various security vulnerabilities for 15,820 IoT backends. The discovery and scanning have been part of a separate research project, in this contribution we focus on the disclosure to the backends' operators in a large-scale coordinated vulnerability disclosure effort, following the latest disclosure guidelines. We discuss what we have learned to inform others who want to engage in large-scale CVD, we compare the steps and tradeoffs of our effort with current CVD suggestions, based on our measurement before and after the disclosure, and we describe how adapting our approach can improve CVD best practices.

1. Introduction

The widespread adoption of the Internet of Things (IoT) has raised a multitude of challenges in ensuring the safety of Information and Communications Technology (ICT) infrastructure and users' privacy on an Internet-wide scale. The number of discovered vulnerabilities has increased rapidly over the past years, many of which affect the growing number of network-connected devices and their, typically publicly reachable, backend services. That is, stakeholders may expose vulnerable network backends and devices that attackers could abuse maliciously. To improve security, including Internet security and IoT device security, vulnerability disclosure has taken an essential role, encouraging active collaboration between security researchers, practitioners, operators, stakeholders, and end-users. For example, to mitigate identified issues before they are exposed and could be abused by attackers. Intuitively, the accuracy of disclosure information shared between the parties determines if security actions are taken promptly and without hesitation, and, with disclosure, the process extends beyond a purely technical understanding of the vulnerability to also include the intricacies and challenges of human-to-human communication.

Coordinated Vulnerability Disclosure (CVD) aims to establish a constructive and practical communication between security researchers and stakeholders, by guiding the researchers in conducting vulnerability disclosure and initially limiting reporting of the discovered vulnerabilities to vendors and affected parties, before eventually publicly revealing them. This allows both sides to collaborate and mitigate the problems before the information becomes public and attackers could abuse it. Finally, the vendor informs the users to act on the discovered and addressed issues. CVD has enhanced vulnerability disclosure, but there are still tradeoffs and concerns to consider [15]. For example, researchers may struggle to locate the correct contact for stakeholders [2], especially when engaging in large-scale notifications. The level of detail to disclose and the required methods to remedy the issue can also trigger unwanted behavior from the stakeholders. Similarly, it can be challenging for the stakeholders to identify the true severity of the disclosed issues, which might hinder their mitigation efforts. Fortunately, prior work proposed well-structured strategies, based on experience from researchers, practitioners, and operators in the last decade, to address these concerns [19, 27]. We conduct our worldwide disclosure process leveraging their suggested methodologies and we provide our own experience as new feedback to improve the CVD process.

In earlier work, we analyzed and measured the security posture of IoT backends, that is, backends speaking IoT-focused protocols (e.g., Message Queuing Telemetry Transport (MQTT)), at scale, focusing on understanding if backends suffer from several known and high impact vulnerabilities [3, 4, 5, 6, 8]. Overall, we identified 15,820 potentially vulnerable backends and prepared the disclosure of the discovered vulnerabilities to the responsible parties. We followed current CVD guidelines to encourage stakeholder cooperation and prompt mitigation. Indeed, our disclosure efforts were successful in that some stakeholders fixed their vulnerable backends. However, we also received unexpected responses and faced difficulties when informing affected parties with complex operational structures. By discussing our experience and actual consideration of each step of our disclosure process, we shed light on formulating an adequate disclosure for large-scale vulnerability disclosure, and we provide new insight on the connection between CVD as a model and implementing it.

1.1. Contributions

In this paper, we make the following contributions:

- We performed large-scale responsible disclosure for IoT backends suffering from vulnerabilities, based

DOI 10.1109/EuroSPW61312.2024.00039

on current Coordinated Vulnerability Disclosure (CVD) best practices, and we measured stakeholders' responses after our disclosure notifications.
- We provide comprehensive guidelines to facilitate large-scale disclosure notifications and handle stakeholder communication.
- We discuss what we learned from our disclosure and we provide practical suggestions to further improve CVD best practices.

2. Background and Related Work

This paper focuses on our Coordinated Vulnerability Disclosure experience. We briefly describe our measurement results. We then discuss recent updates of CVD practices and how we adopt them into our disclosure operation.

2.1. Security Assessment

We leveraged Shodan [22] to identify backends that speak common IoT communication protocols, like MQTT, Constrained Application Protocol (CoAP), and Extensible Messaging and Presence Protocol (XMPP). Together with the IP addresses and hostnames for the backends, we also collected Connection Codes and geolocation information. We then performed our security evaluation, focusing on three main attack classes: information leakage, weak authentication, and denial of service. We define the testing pipeline for each protocol and provide more details in our prior work [25]. In this paper, we instead focus on the disclosure process for the discovered vulnerable backends. Overall, we discovered numerous backends suffering from known vulnerabilities, identified via Common Vulnerability Exposures (CVEs), and other common security pitfalls, such as not requiring authentication or unintended exposed access. We then compiled a list of vulnerable backends with their IP addresses, CVEs or vulnerabilities if no CVE was assigned, and other issues that we identified, which we utilized to disclose our findings to the backend operators following current CVD best practices.

We first analyzed the security of the backends in September 2022, and later repeated our analysis in September 2023. Table 1 shows the number of vulnerable MQTT backend for September 2023 for each of the five CVEs that we tested for. Since CVE-2018-12550 and CVE-2018-12551 always occurred together, we combine them in the table. All CVEs were published over three years before our scan and mitigations have been available, in the form of an updated version since the vulnerability become public, which should be ample opportunity to update and resolve the issue. We do not provide data for CoAP and XMPP because our disclosure process for backends speaking these protocols is still ongoing.

2.2. Coordinated Vulnerability Disclosure

CVD is a well-known process that has become accepted best practice and the standard operating procedure for most security researchers. Its idea is to contact vendors or operators of affected services and products, like backends, before publicly disclosing the existence of vulnerabilities, so that they can address the vulnerability during an embargo period. There is no minimum or maximum time length for an embargo, but approximately 90 days has become typical. Only after the embargo period, the parties publish about the vulnerability, so that others can take knowledge about the vulnerability into account for their decision-making processes, including the remedy, or learn from it.

Table 1. We report the number of backends affected by CVEs we tested for grouped by protocol. CVE-2017-7655, CVE-2018-12550, and CVE-2018-12550 affect the Eclipse Mosquitto MQTT server, CVE-2018-19417 affects Contiki-NG, and CVE-2019-9749 affects Fluent Bit.

Protocol	CVE	CVSSv3	# Found
MQTT	CVE-2018-12550 [4] CVE-2018-12551 [5]	8.1	6,236
	CVE-2017-7655 [3]	7.5	4,241
	CVE-2018-19417 [6]	10.0	186
	CVE-2019-9749 [8]	7.5	151

This basic approach is widely accepted and described [16, 17, 26]. In practice, however, this turns out to be more complicated, especially if multiple parties are involved [10], and this procedure must be adapted. This has led the University of Twente, where we initiated our disclosure process, to define an outgoing vulnerability disclosure policy [13, 21]. This policy also explains to the recipients of our notifications how we approach the disclosure process, for example, clarifying our timelines and expectations. The policy builds on the disclosure guidelines of the Dutch National Cyber Security Centre (NCSC-NL) [20]. For our CVD, we strictly follow these policies.

We also apply methods recommended by the Dutch Institute for Vulnerability Disclosure (DIVD) and will describe our results and experiences using their framework and C-HIP model [28]. Moreover, we utilize the investigation of the state of CVD by van der Horst to address difficulties of large-scale disclosure efforts for our process [14]. The DIVD has extensive experience in large-scale vulnerability disclosures, but has focused primarily on disclosing critical corporate infrastructure bugs, such as the Citrix vulnerability in early 2020 [7]. In contrast, we investigate the disclosure of IoT backend vulnerabilities, a setting that is distinctly different to critical corporate infrastructure in terms of security considerations.

We locate the contacts of IoT backends via the WHOIS databases [18, 24]. We use simple language for our disclosure notifications, see Appendix A, and we also set up a static webpage, detailing the research to inform stakeholders[1], akin to prior work [1, 12]. Since some IoT backends are hosted on public clouds, it can be challenging to identify the actual operator. Therefore, we also notified cloud providers, with the request to forward our notification to the affected customer. During our interactions with stakeholders, we observed that not all explicitly specify their disclosure policies or procedures. Unfortunately, the recommended practice for vendors or operators to serve an informative `security.txt` file containing security contact-related information for security researchers to contact [11] is not yet widely used. Moreover, it only applies to web-based systems or requires knowledge about who the actual operator of the vulnerable service is and their website, which is particularly difficult to determine for services

1. https://iot-disclosure2023.dacs.utwente.nl/

deployed on public clouds, further limiting its usefulness for our work.

3. Methodology

Following, we describe how we structured our large-scale disclosure. We first gathered the required contact information and composed the disclosure notifications to the corresponding contacts. We then notified the stakeholders at scale and tracked the conversations via ticketing methods.

3.1. Disclosure Procedures

3.1.1. Contact Information Retrieval. We identified 15,820 IP addresses of vulnerable IoT backends based on our security assessment [25]. We retrieved the corresponding contact information for the IP addresses via the WHOIS databases of APNIC, RIPE, etc. Eventually, we gathered email addresses and contact details for 15,046 IP addresses (95.11%). Unfortunately, for some backends, the WHOIS database contains no useful information, that is, it simply returns a general abuse contact for the corresponding Regional Internet Registry.

3.1.2. Disclosure Messages. We composed our disclosure notification template to ensure our disclosure messages are clear and informative to bootstrap conversations with the stakeholders (see Listing 1). In our initial message, we inform each stakeholder about the identified vulnerable backends, providing information on the IP addresses, CVEs, and vulnerabilities we found. We briefly describe the severity of the vulnerabilities that require the stakeholder's prompt action and we point them to a dedicated website with more detailed information. The website provides an overview of the issues, links to additional material by MITRE, and an email address to contact us.

We reply with a second message (see Listing 2) when stakeholders return to us with additional questions. Our reply includes details of the vulnerable backends and suggestions for mitigations. The exact content may vary depending on the security issues of backends and the stakeholders' responses.

3.2. Notification System

To notify the contacts at scale, we set up a dedicated email account with the domain name utwente.nl to send our disclosure notifications. We chose this approach to ensure our notifications have a high chance of being delivered [14, 23] and that they are not discarded by recipients, for example, because of a low sender reputation.

We used the email server of the University of Twente, which uses Microsoft Exchange. We sent out notifications in batches, due to limits enforced by Microsoft Exchange service policies. To facilitate large-scale disclosure, we implemented a custom notification script in Python, which interacted with our SMTP server to send disclosure notifications while respecting the limits of the mail server. For each message, our implementation also logs the timestamp, contact email addresses, and identifiers.

Before sending notifications, we grouped the vulnerable backends by contact address, that is, we combined those for which the contact address is the same. This grouping enabled us to drastically reduce the emails we had to send from 15,046 to 2,132. On average, we reported seven vulnerable backends per disclosure notification (max 1,555). We ensure brevity and understandability of our disclosure message by listing only a subset of up to ten vulnerable backends in the body of our message and including the complete list of vulnerable backends as a text file attachment.

Table 2. We manually cluster the stakeholder responses in four main categories. We mark as *Automatic Response* those that were automatically sent by ticket systems etc. We flag as *Manual* the responses of interested parties that contacted us after our disclosure emails, split into *Question* for inquiries about more details and *Update* for informing us about the issue being resolved. We denote as *Delivery Failure* notifications that we sent but which were returned by the recipient email server as undeliverable.

Type		Core Message	Count
Automatic Response		Thank you for your disclosure. We received your message and will inform the responsible parties.	428
Manual	Question	Please provide more details about the vulnerabilities and remedies.	32
	Update	We mitigated the disclosed issues.	13
Delivery Failure		Contact not reachable.	700

4. Results

Following, we discuss the results of our disclosure effort with the stakeholders, how we initiated the conversations, the challenges we encountered, and what we learned.

4.1. Measurements and CVD Timeline

We performed our initial security assessment in September 2022, followed by a second assessment in September 2023, to understand how the security posture of backends evolved. After evaluating the results of both measurements, we sent disclosure notifications from November 10th, 2023, to November 22nd, 2023. Overall, we sent 2,132 emails to notify contacts associated with 15,046 IP addresses. We received 32 requests for additional information. After responding to those 32 requests, 28 stakeholders then initiated an active conversation with us, who we helped locate and fix their issues. The remaining 4 did not further respond and did not fix the identified security issues. An additional 311 stakeholders fixed the identified vulnerabilities without responding to us.

4.2. CVD with Stakeholders

After disclosing the issues, we received 1,173 responses to our notifications from stakeholders between November 2023 and January 2024. Table 2 summarizes the numbers. Most of them (428, 36.48%) were automatic responses from the stakeholders' ticketing systems. We also received replies from individuals, such as service operators, company incident response teams, or users, who were notified by the cloud service providers. Unfortunately, we received 700 (59.67%) delivery failure notifications, which we discuss further in Section 4.2.3.

4.2.1. Automatic Responses. Automatic responses were generated by ticketing systems. Most stakeholders, such as cloud service providers, universities, and business organizations, use ticketing systems to handle incoming email messages or abuse reports. These systems received our disclosure notifications and forwarded them to the responsible parties within these entities. In 27 cases, the ticketing systems required us to sign in to their portal with automatically generated accounts to confirm our disclosure, likely to prevent spam. To ensure our notifications were delivered successfully, we followed the corresponding instructions to confirm our messages. However, three backends not only required us to provide our contact information but also requested us to accept their policies, for example, privacy policies, which were only available in their native languages. We opted to not accept these policies with uncertain impact, which may have led these parties to ignore our disclosures. Besides the aforementioned situations, 51 ticketing systems correctly forwarded our notifications to the affected parties. Seven cloud providers' systems informed us about updates on their communication and progress with their customers. In two cases, they established direct communication between their customers and us through their ticketing system.

4.2.2. Manual Responses. We consider replies to our disclosure notifications that are from stakeholders who are directly responsible for the vulnerable backends as manual responses. We separate them into two classes, first, questions concerning the disclosed vulnerabilities, and, second, updates or comments by them regarding the coordinated disclosure process (e.g., that a vulnerability was mitigated). We received 28 positive replies and valuable suggestions, as well as 4 complaints. In Section 4.3, we discuss how we handled all manual responses and the corresponding active coordination of our CVD.

4.2.3. Delivery Failures. Not all our disclosure notifications were delivered successfully to the responsible parties. Delivery failures can occur for several reasons, such as contact addresses no longer being valid or the recipient email server employing restrictive filtering rules. We received 700 delivery failure notifications and attempted to cross-check these with our outgoing emails, as our disclosure notifications could have multiple recipients, a 1-to-1 matching is not possible. For 103 out of 2,132 messages (<5%), we received failure notifications for all recipients. This affected 210 backends (<2%) for which we were not able to notify the operator at all. For the remainder, notifications failed to deliver only for some recipients, but not all, and we assume that at least one notification was successfully delivered. We describe some reasons for delivery failures below but have not been able to categorize all cases exhaustively, which is why we can not provide exact numbers.

If a delivery failure occurred, we could not notify the affected parties if we could not find any alternative contact method. The contact email addresses' filtering rules might have also blocked our messages. Some email servers restrict the length of content messages, attachment type, and file size, including to abuse and other contact addresses for operators. To our surprise, some even reject small text file attachments. This was problematic for our disclosure because we included a text file attachment with the complete list of affected vulnerable IP addresses, ports, timestamps of scanning, and CVEs for all emails and only included up to 10 results in the email itself to prevent messages being marked as spam. If the recipient server rejects our text file attachments, then this further complicates the process of notifying the affected stakeholders.

4.3. Handling the Responses from Stakeholders

A core component of CVD is establishing a trusted and honest communication channel with the stakeholders. The goal of our initial disclosure message is to draw the recipient's attention to the discovered issues and establish such a communication channel, to coordinate with the affected stakeholder on the embargo timeline and how to remedy the reported issues. During this process, the stakeholders sent us questions on the details of our findings, suggestions on our disclosure method, and different perspectives on our vulnerability findings. We also received complaints regarding our disclosure operations. Following, we detail in depth how we interact with the stakeholders, what we have learned from the disclosure process, and how we incorporated their suggestions into our later disclosures.

In our initial batch of notifications, we sent notifications to recipients associated with vulnerable IP addresses for MQTT backends, both with and without detected CVEs. We also included additional information and a reference to our website that included more information about our scans and disclosure practices.

After our notifications, 32 recipients replied to our message requesting additional information. We provided them with additional information (see Listing 2), such as how we found the vulnerability and security advice on mitigating the issues. After this, 28 respondents reported that they were working on locating the problems and informing related parties, such as their clients and other operators.

Over 60% of 15,820 backends were hosted on cloud platforms, preventing us from directly notifying the operator and only letting us notify the cloud providers. While we could not contact the operators directly, the cloud providers relayed our notifications to their customers and other responsible parties. After the responsible parties were informed, we received their feedback by email directly or through the cloud providers' ticketing systems.

We also received eight complaints that the vulnerable backends were not part of the contacts' operation and not under their control. Since we collected the contact email addresses via WHOIS, the contact addresses may have been inaccurate.

4.4. Before and After Our Disclosures

We conducted a third security assessment in January 2024, after our disclosures, to measure and analyze if the stakeholders mitigated the vulnerabilities and adopted our recommendations to improve their security.

Figure 1 provides an overview of the results, as a percentage of the 14,836 backends that we successfully notified (notifications for 15,046 backends total, but delivery failed for all contacts for 210 backends, which we consider as not notified). The operators of 52 (0.35%) backends addressed

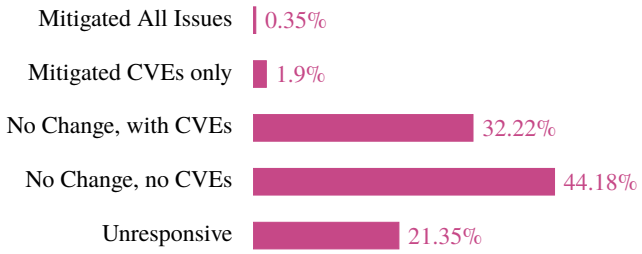

Figure 1. Security posture of vulnerable IoT backends as of January 2024, approximately two months after we sent our disclosure emails in November 2023. All numbers are relative to the 14,836 IoT backends for which at least one disclosure notification was delivered, that is, the number of backends we consider as successfully notified.

the vulnerabilities and also implemented our security recommendations, such as requiring authentication. Meanwhile, for 282 (1.90%) backends the operators only mitigated the reported vulnerabilities identified by their CVEs, but they did not implement our other recommendations. For example, we could still observe exchanged messages, without needing to authenticate. Unfortunately, 3,168 (21.35%) backends became unresponsive, likely because of IP address instability. Although only few 334 (2.25%) backends mitigated vulnerabilities, this is not unexpected given the limited accuracy of WHOIS contact addresses. We have received a reply from 31 service providers that they have forwarded our disclosure notifications to their customers, but we cannot be certain that they reached the customers unless they replied to us.

Worse, 4,780 backends (32.22%) remain vulnerable. The remaining 6,554 (44.18%) backends have no known vulnerabilities tracked by CVEs, but they are publicly exposing resources to the Internet and do not require any authentication, which, for example, leaves broadcast messages unprotected. From our CVD, we learned that these backends range from university projects to commercial applications to enterprise-grade IoT backend services.

5. Ethical Considerations

Throughout our measurements [25] and the following disclosure process, we kept the impact of our work in mind and followed best practice and CVD guidelines. We received approval for our measurements from the Ethics Review Board (ERB) at the University of Twente.

In this paper, we share our CVD approach and experience, so that the community and others can benefit from them and further improve CVD processes. Our disclosure operation was the first large-scale CVD implementing the outgoing vulnerability disclosure policy of the University of Twente [13, 21]. No respondents remarked on this policy explicitly, but it provided a public guideline for timelines and our predictable actions. Some respondents might have decided to not respond and simply mitigate the issues without engaging with us after taking note of it. We recommend other researchers also clearly communicate and publish such an outgoing CVD policy.

6. Reflections on our CVD

We set up our disclosure notification operation based on the guidelines given by DIVD [14], which is why we compare it to the DIVD suggestions, the existing theories, and best practices. We follow the Communication-Human Information Processing (C-HIP) model [28], focusing on the following elements, to better present the comparison.

6.1. Channel & Delivery

We used email as our notification method. The vulnerable backends are maintained by diverse organizations with different disclosure practices across countries. Correspondingly, for our large-scale disclosure notification for these backends, the WHOIS abuse contact information for the affected IP addresses is the most reliable way to reach as many responsible parties as possible. It was not feasible to establish a second channel, such as a notification-sharing platform, since the contacts are not within the scope of an organization. For backends hosted on cloud infrastructures, we informed the providers and asked them to forward our messages. We received copies for some cloud providers, who sent a request to their clients to fix the issues or take down the backends. For other disclosures to cloud providers, we were not privy to any notification to their customers.

6.2. Attention Maintenance & Comprehension

The recipients of disclosures must quickly identify them and keep track of the vulnerability disclosure process. In our disclosure template (see Appendix A), we provided brief but essential information to identify the vulnerable backends. Our goal was to draw the recipients' attention without overwhelming them with information. However, there is a trade-off between message length and effectiveness. Stakeholders typically use email ticketing systems, which record our disclosures and inform the system operators. Certain cloud providers have also stringent restrictions on email content, filtering messages that are too long or include attachments that are too large. Using HTML-based content, such as colors, different fonts, or other graphic content in our disclosure message is also not ideal, because it could trigger spam filters.

6.3. Attitudes & Motivation

To communicate positively and encouraging with the affected stakeholders and find a balance to motivate them to actively engage in the disclosure process, we have incorporated stakeholders' responses and suggestions into our notification methods. Early respondents requested that we include a brief description of the severity of vulnerabilities, IP addresses, port numbers, and preferably corresponding CVEs, which they see as the optimal information to include in the disclosure messages. We remedy the fact that CVEs may not be understood by all recipients equally by providing additional security suggestions in simple text, which we hope helps reduce the efforts of stakeholders to mitigate the vulnerabilities. We also specify our contact information in our disclosure email, and we provide a link to our webpage that provides more details. Contacts that actively engaged in the CVD informed us that they found this helpful. Another important concern is the language to communicate with the stakeholders. Since our disclosure

were large-scale and world-wide, we composed our messages in English. Interestingly, we received 112 responses, mainly from stakeholders' ticketing systems, that were in the stakeholders' native languages, which increased our efforts to react to those responses.

6.4. Receiver & Behavior

WHOIS email contact information may not always be accurate, which means that we could potentially leak information on vulnerabilities to the wrong recipient. Indeed, this was the case for eight recipients, who complained that we informed the wrong parties. We aim to mitigate this situation by limiting the vulnerability details in our first message and not providing every detail. However, this can conflict with the interest of stakeholders, who prefer a clear and informative initial message, for example, to triage it internally. Another consideration is the tendency of stakeholders to start mitigating the vulnerabilities without notifying us. Whether stakeholders informed us about their process appears to also depend on their own vulnerability management policies and willingness to participate in our CVD. Another contact method than email could improve on these aspects, for which our next step is to construct a reliable disclosure channel with the stakeholders. The responses we received show that well-known cloud providers and organizations host numerous IoT backends. Establishing long-term connections and communication channels with these organizations could allow us to directly reach the responsible parties and improve CVD effectiveness.

7. Future Work

Traditionally, meta data regarding IP addresses, like contact addresses, is accessible through WHOIS, which is unstructured text for which it can be difficult to retrieve the correct email address. The Registration Data Access Protocol (RDAP),[2] which provides structured access to this meta data, has been introduced as a successor to WHOIS and could improve delivery reliability for CVD notifications. Unfortunately, while RDAP is in the process of being deployed, WHOIS remains more widely supported, reliable, and more frequently updated [9]. Future work, especially once RDAP becomes more widely used, should explore the benefits of RDAP for retrieving accurate contact information in more depth. Moreover, given its extensibility, RDAP provides an opportunity to include a disclosure contact address, which could improve CVD effectiveness.

8. Conclusion

Coordinated Vulnerability Disclosure (CVD) is a well-established and proven approach to unite security researchers, vendors, operators, and end users to mitigate the identified vulnerabilities in a joint effort. Unfortunately, as we have experienced in our large-scale CVD of vulnerabilities spanning over 15,820 IoT backends that we discuss in this paper, it remains challenging for researchers to set up reliable infrastructure to facilitate large-scale CVD and establish practical and trusted communication with stakeholders when considering the intricacies of real-world deployments. Following current best CVD practices to contact operators, we could still not inform the operators of 984 backends (6.3%) about their vulnerable systems. CVD does lead to improvements in security for IoT backends, these are limited to a low 2.28% of backends, while a total of 77.5% of backends see no security improvements after disclosure. That is, our findings suggest a troubling immaturity regarding backend security awareness and the adoption of CVD in the IoT ecosystem. Overall, leveraging our own experience, in this paper, we provide new insight on the struggles when performing large-scale CVD at the example of IoT backends and we make new suggestions to improve current CVD practices.

2. https://about.rdap.org/

Acknowledgements

This work is based on research supported by the INTERSCT project, Grant No. NWA 1160.18.301, funded by the Netherlands Organisation for Scientific Research (NWO), the Vienna Science and Technology Fund (WWTF) and the City of Vienna [Grant ID: 10.47379/ICT19056], the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972, SBA Research (SBA-K1), a COMET center funded by BMK, BMDW, and the state of Vienna, and the Internet Society Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the respective funding agencies.

We appreciate and are grateful for the support from Andrea Continella who connected the authors and helped prepare the disclosure process, the email infrastructure team, LISA, at the University of Twente, as well as all stakeholders who actively engaged in our CVD effort and provided valuable insights for our disclosure process.

References

[1] K. Borgolte, S. Hao, T. Fiebig, and G. Vigna. "Enumerating Active IPv6 Hosts for Large-scale Security Scans via DNSSEC-signed Reverse Zones." In: *Proceedings of the 39th IEEE Symposium on Security & Privacy (S&P)*. May 2018. DOI: 10.1109/SP.2018.00027.

[2] O. Çetin, C. Gañán, M. Korczyński, and M. van Eeten. "Make Notifications Great Again: Learning How to Notify in the Age of Large-Scale Vulnerability Scanning." In: *Proceedings of the 16th Workshop on the Economics of Information Security (WEIS)*. June 2017. URL: https://infosecon.net/workshop/downloads/2017/pdf/Make_Notifications_Great_Again:_Learning_How_to_Notify_in_the_Age_of_Large-Scale_Vulnerability_Scanning.pdf (visited on 05/19/2024).

[3] *CVE-2017-7655*. Apr. 11, 2017. URL: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2017-7655 (visited on 05/19/2024).

[4] *CVE-2018-12550*. June 18, 2018. URL: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-12550 (visited on 05/19/2024).

[5] *CVE-2018-12551*. June 18, 2018. URL: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-12551 (visited on 05/19/2024).

[6] *CVE-2018-19417*. Nov. 21, 2018. URL: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-19417 (visited on 05/19/2024).

[7] *CVE-2019-19781*. Dec. 13, 2019. URL: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-19781 (visited on 05/19/2024).

[8] *CVE-2019-9749*. Mar. 13, 2019. URL: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-9749 (visited on 05/19/2024).

[9] S. Fernandez, O. Hureau, A. Duda, and M. Korczynski. "WHOIS Right? An Analysis of WHOIS and RDAP Consistency." In: *Proceedings of the 16th Passive and Active Measurement Conference (PAM)*. Springer, Mar. 2024. DOI: 10.1007/978-3-031-56249-5_9.

[10] Forum of Incident Response and Security Teams (FIRST). *Guidelines and Practices for Multi-Party Vulnerability Coordination and Disclosure*. June 19, 2020. URL: https://www.first.org/global/sigs/vulnerability-coordination/multiparty/guidelines-v1.1 (visited on 05/19/2024).

[11] E. Foudil and Y. Shafranovich. *A File Format to Aid in Security Vulnerability Disclosure*. Tech. rep. 9116. Request for Comments (RFC), Apr. 2022. DOI: 10.17487/RFC9116.

[12] J. van der Ham and R. van Rijswijk-Deij. "Ethics and Internet Measurements." In: *Journal of Cyber Security and Mobility, Volume 5, Issue 4* (Oct. 2017). URL: http://dx.doi.org/10.13052/jcsm2245-1439.543.

[13] J. van der Ham, A. Continella, P. de Willigen, and D. Reidsma. *University of Twente Policy for Coordinated Vulnerability Disclosure in Research*. Mar. 22, 2023. URL: https://www.utwente.nl/en/service-portal/research-support/procedures-facilities/coordinated-vulnerability-disclosure-policy-for-research (visited on 05/19/2024).

[14] M. van der Horst. *Global Vulnerability Vigilance: Timely Disaster Notification using Internet-Scale Coordinated Vulnerability Disclosure*. July 11, 2023. URL: https://scripties.uba.uva.nl/search?id=record_54279 (visited on 05/19/2024).

[15] A. D. Householder and J. Spring. "Are We Skillful or Just Lucky? Interpreting the Possible Histories of Vulnerability Disclosures." In: *Digital Threats: Research and Practice* 3.4 (Feb. 7, 2022). DOI: 10.1145/3477431.

[16] ISO/IEC JTC 1/SC 27. *ISO/IEC 29147:2018 Information technology – Security techniques – Vulnerability disclosure*. Tech. rep. 29147. International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC), Oct. 2018. URL: https://www.iso.org/standard/72311.html (visited on 05/19/2024).

[17] ISO/IEC JTC 1/SC 27. *ISO/IEC 30111:2019 Information technology - Security techniques - Vulnerability disclosure*. Tech. rep. 30111. International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC), Oct. 2019. URL: https://www.iso.org/standard/69725.html (visited on 05/19/2024).

[18] F. Li, Z. Durumeric, J. Czyz, M. Karami, M. Bailey, D. McCoy, S. Savage, and V. Paxson. "You've Got Vulnerability: Exploring Effective Vulnerability Notifications." In: *Proceedings of the 25th USENIX Security Symposium (USENIX Security)*. USENIX Association, Aug. 2016. URL: https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/li (visited on 05/19/2024).

[19] A. Nakajima, T. Watanabe, E. Shioji, M. Akiyama, and M. Woo. "A Pilot Study on Consumer IoT Device Vulnerability Disclosure and Patch Release in Japan and the United States." In: *Proceedings of the 14th ACM Asia Conference on Computer and Communications Security (ASIACCS)*. Association for Computing Machinery (ACM), July 2, 2019. DOI: 10.1145/3321705.3329849.

[20] National Cyber Security Centre of The Netherlands (NCSC-NL). *Coordinated Vulnerability Disclosure: the Guideline*. Oct. 2, 2018. URL: https://english.ncsc.nl/publications/publications/2019/juni/01/coordinated-vulnerability-disclosure-the-guideline (visited on 05/19/2024).

[21] D. Reidsma, J. van der Ham, and A. Continella. "Operationalizing Cybersecurity Research Ethics Review: From Principles and Guidelines to Practice." In: *Proceedings of the 2nd Workshop on Ethics in Computer Security (EthiCS)*. Internet Society, Feb. 2023. DOI: 10.14722/ethics.2023.237352.

[22] *Shodan*. 2023. URL: https://www.shodan.io (visited on 10/14/2023).

[23] B. Stock, G. Pellegrino, F. Li, M. Backes, and C. Rossow. "Didn't You Hear Me? - Towards More Successful Web Vulnerability Notifications." In: *Proceedings of the 25th Network and Distributed System Security Symposium (NDSS)*. Internet Society, Feb. 2018. DOI: 10.14722/ndss.2018.23171.

[24] B. Stock, G. Pellegrino, C. Rossow, M. Johns, and M. Backes. "Hey, You Have a Problem: On the Feasibility of Large-Scale Web Vulnerability Notification." In: *Proceedings of the 25th USENIX Security Symposium (USENIX Security)*. USENIX Association, Aug. 2016. URL: https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/stock (visited on 05/19/2024).

[25] C. Tagliaro, M. Komsic, A. Continella, K. Borgolte, and M. Lindorfer. "Large-Scale Security Analysis of Real-World Backend Deployments Speaking IoT-Focused Protocols." May 2024. arXiv: 2405.09662 [cs.CR].

[26] "Understanding the digital security of products." In: *OECD Digital Economy Papers* (Feb. 9, 2021). DOI: 10.1787/abea0b69-en.

[27] T. Walshe and A. C. Simpson. "Coordinated Vulnerability Disclosure programme effectiveness: Issues and recommendations." In: *Computers & Security* 123 (Dec. 2022). DOI: 10.1016/j.cose.2022.102936.

[28] M. S. Wogalter, D. M. DeJoy, and K. R. Laughery. “Organizing Theoretical Framework: A Consolidated Communication-Human Information Processing (C-HIP) Model.” In: *Warnings and Risk Communication*. CRC Press, 1999. ISBN: 9780429224867.

A. Disclosure Email

Dear Sir/Madam,

As a collaborative project from the University of Twente and TU Wien, we have been investigating the security of publicly available IoT services. We examined the server endpoints that IoT devices and smartphone apps connect to. Our study has revealed a security issue that might require your immediate attention. We conducted a scan and found the following vulnerable machines (IP address, Found Time, and CVEs):

192.0.2.42 2022-08-24 00:03:02.000000000 +0000 CVE-2018-12550, CVE-2018-12551

You can find out more about our study and an explanation of the results on our project website: https://iot-disclosure2023.dacs.utwente.nl

These findings have been kept confidential and all the rules stated by our Dutch National Cyber Security Center (NCSC) [1] were adhered to during this research. We are planning to follow the procedure as outlined in our public disclosure policy [2].

If you are not running this server yourself, but know the responsible party (e.g., because they are your customer), please forward this information to them, you have our explicit approval for this.

Regards,
IoT Disclosure Project, UTwente

[1]: https://english.ncsc.nl/get-to-work/implement-a-cvd-policy/finding-vulnerabilities-in-it-systems
[2]: https://www.utwente.nl/en/service-portal/research-support/procedures-facilities/coordinated-vulnerability-disclosure-policy-for-research

Listing 1: Initial Disclosure Email

Dear Sir/Madam,

As a collaborative project from the University of Twente and TU Wien, we have been investigating the security of publicly available IoT backends. We examined the server endpoints that IoT devices and smartphone apps connect to. Our study has revealed a security issue that might require your immediate attention. We conducted a scan and found the following vulnerable machine (IP address, Found Time, CVEs):

192.0.2.42 2022-08-24 00:03:02.000000000 +0000 CVE-2018-12550, CVE-2018-12551

In our scanning method, we attempted to subscribe to your MQTT services with wildcards "$SYS/#" and "#" on Port 1883. We collected the number of topics (680) and the number of connected clients (2) at the time of our scan. We did not explore your system further. Yet we believe an attacker can use the same way to extract all MQTT content and may get access to sensitive information such as Personal Identifiable Information or security device status of associated systems.

Our suggested improvements are:

* Adopt some authentication measures, e.g., password-based authentication.
* Adopt some Access Control Lists to prevent all users from reading all messages, e.g., only admins can read certain sensitive topics.
* Encrypt the communication, e.g., by using TLS (some lightweight versions of it can work with IoT devices, e.g., https://www.wolfssl.com/)
* If the broker does not have to be exposed to the entire Internet, protect it either behind a firewall (blocking incoming connections from outside the organization) or a NAT.

You can find out more about our study and an explanation of the results on our project website: https://iot-disclosure2023.dacs.utwente.nl

These findings have been kept confidential and all the rules stated by our Dutch National Cyber Security Center (NCSC) [1] were adhered to during this research. We are planning to follow the procedure as outlined in our public disclosure policy [2].

If you are not running this server yourself, but know the responsible party (e.g., because they are your customers), please forward this information to them, you have our explicit approval for this.

Regards,
IoT Disclosure Project, UTwente

[1]: https://english.ncsc.nl/get-to-work/implement-a-cvd-policy/finding-vulnerabilities-in-it-systems
[2]: https://www.utwente.nl/en/service-portal/research-support/procedures-facilities/coordinated-vulnerability-disclosure-policy-for-research

Listing 2: Supplemental Followup Email