

Catch me if you can

Thomas, Richard J.; Gardiner, Joseph; Chothia, Tom; Samanis, Emmanouil; Perrett, Joshua; Rashid, Awais

DOI:

[10.1145/3411498.3419970](https://doi.org/10.1145/3411498.3419970)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Thomas, RJ, Gardiner, J, Chothia, T, Samanis, E, Perrett, J & Rashid, A 2020, Catch me if you can: an in-depth study of CVE discovery time and inconsistencies for managing risks in critical infrastructures. in *CPSIoTSEC'20: Proceedings of the 2020 Joint Workshop on CPS&IoT Security and Privacy*. Association for Computing Machinery (ACM), pp. 49-60, CPSIoTSec: The Joint Workshop on CPS & IoT Security and Privacy, 9/11/20. <https://doi.org/10.1145/3411498.3419970>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Catch Me If You Can: An In-Depth Study of CVE Discovery Time and Inconsistencies for Managing Risks in Critical Infrastructures

Richard J. Thomas
University of Birmingham
Birmingham, UK
r.j.thomas@cs.bham.ac.uk

Joseph Gardiner
Bristol Cyber Security Group,
University of Bristol
Bristol, UK
joe.gardiner@bristol.ac.uk

Tom Chothia
University of Birmingham
Birmingham, UK
t.p.chothia@cs.bham.ac.uk

Emmanouil Samanis
Bristol Cyber Security Group,
University of Bristol
Bristol, UK
manolis.samanis@bristol.ac.uk

Joshua Perrett
Bristol Cyber Security Group,
University of Bristol
Bristol, UK
wx19297@bristol.ac.uk

Awais Rashid
Bristol Cyber Security Group,
University of Bristol
Bristol, UK
awais.rashid@bristol.ac.uk

ABSTRACT

Industrial Control Systems (ICS) are central to the operation of critical national infrastructure (CNI) such as oil and gas, water treatment, power generation and transport systems. Effective risk management to mitigate large-scale disruption to societies and economies depends on both timely information about vulnerabilities and the consistency of this information. The longer the vulnerabilities remain “in the wild” or a lack of consistency in vulnerability reporting, the greater the impact on CNI operators’ ability to systematically understand and mitigate the risks. In this paper, we focus on vulnerabilities identified and reported in Siemens ICS devices, which hold the largest share of the market. We undertake an in-depth analysis of 207 CVEs, identifying the time over which vulnerabilities were ‘in the wild’ before being discovered and advisories issued, and examine issues with the correctness of CVE information. We find that, on average, a vulnerability is ‘in the wild’ for 5.3 years, and that many CVEs do not correctly reflect and state the affected devices as Common Platform Enumerations (CPEs). Based on our findings, we propose a set of guidelines to improve the reporting and consistency of ICS CVE information.

CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyber-physical systems; Maintainability and maintenance**; • **Security and privacy** → **Embedded systems security**.

KEYWORDS

Industrial Control Systems; ICS Security; SCADA; Operational Technology, OT; CPS; Cyber Security; Vulnerabilities

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CPSIoTSEC’20, November 9, 2020, Virtual Event, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8087-4/20/11...\$15.00
<https://doi.org/10.1145/3411498.3419970>

ACM Reference Format:

Richard J. Thomas, Joseph Gardiner, Tom Chothia, Emmanouil Samanis, Joshua Perrett, and Awais Rashid. 2020. Catch Me If You Can: An In-Depth Study of CVE Discovery Time and Inconsistencies for Managing Risks in Critical Infrastructures. In *2020 Joint Workshop on CPS&IoT Security and Privacy (CPSIoTSEC’20)*, November 9, 2020, Virtual Event, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3411498.3419970>

1 INTRODUCTION

A vulnerability within an ICS device can have a catastrophic effect if discovered and subsequently exploited by an attacker, particularly if that device is used within critical national infrastructure (CNI). Within the European Union and United Kingdom, the Network and Information Systems (NIS) Directive placed responsibility of security upon asset owners [22]. This shift of not only protecting the traditional corporate IT environment but also operational technology (OT)¹ systems presents a number of challenges. OT systems typically remain unchanged for a number of decades after deployment, compared to typical IT refresh cycles of at most 5 years. Safety and management of process is pivotal within OT, whereas IT systems concern themselves with business continuity, a well-understood field with standardised security practices. In OT, this standardisation is fairly recent, with OT Security in scope of IEC 62443. Additionally, supply chain security has an important role, where asset owners require accurate and actionable information to manage risk to their environments, implement effective upgrade programs and continuously improve their security. Any vulnerability reported which could apply to an asset owner’s infrastructure must therefore contain actionable information at the right time. If there is any inconsistency or inaccuracy in that information, vulnerabilities may remain unaddressed in that infrastructure with potentially serious consequences if exploited.

Whilst attention to ICS security has increased since the discovery of Stuxnet, more vulnerabilities, in particular those existing in legacy devices, are being discovered [9]. In order to understand the risk that exists to an asset owner’s infrastructure, ICS vulnerability

¹Operational Technology comprises of systems involved in plant and process automation.

Vendor	CVE Count
Siemens	424
Schneider Electric	167
Advantech	140
Moxa	91
Rockwell Automation (inc. Allen-Bradley)	76

Table 1: Top 5 Vendors by ICS CVE count.

reports, including CVEs (Common Vulnerabilities and Exposures), are primary sources of information, driving decisions on how to react. These, however, lack vital details around how long the vulnerability was “in the wild” before being disclosed by the vendor (our first research question), and, thus the window of exposure for the asset owner. By having this information, the asset owner has a clearer understanding of these risks. They can then identify whether unexpected behaviour exhibited in their infrastructure occurred during this window and understand the potential impact to their environment. Without understanding their potential exposure, or being able to confidently state that they are not affected, the risk of exploitation is not appreciated and asset owners may overlook the risk. This is precisely the case for Triton (CVE-2018-7522/8872), where safety systems were unexpectedly failing-safe, prompting the asset owner to investigate, leading to the discovery of malware targeting the system [11]. Additionally, inconsistencies in the CVE description could have a further impact. Therefore, timeliness and accuracy is vital, informing asset owners how to plan, mitigate and understand their security posture. If this information is flawed in any way, it is not possible to act and protect their infrastructure².

Due to the amount of manual work required to perform this analysis, we chose to review vulnerabilities from a single vendor for this study – Siemens. This was for a number of reasons – Siemens is one of the largest ICS vendors in the market [3], with prominence in the production of industrial control and automation systems and the ICS CVE prevalence available for analysis. Of all ICS vendors with reported vulnerabilities, Siemens has the highest number of assigned CVEs, providing a large dataset from one manufacturer. Rockwell, another dominant vendor, does not have the volume of CVEs compared to Siemens, which would otherwise impact the effectiveness of our analysis. It is important to note, however, that whilst Siemens has highest CVE prevalence for ICS-related vulnerabilities (the Top 5 are shown in Table 1), this should not be interpreted that they are considered more vulnerable than another. This is partly due to dominance in the market, and their increased focus on improving the security of their products, and, subsequently reporting vulnerabilities. After the discovery of Stuxnet in 2010, Siemens have put a far greater effort into their ProductCERT team and vulnerability disclosure process [19], which we expect should provide a more efficient vulnerability disclosure process meaning CVEs are assigned sooner after a vulnerability is disclosed.

Our dataset is established from CVEs listed in US-CERT ICS Advisories, which tracks ICS-related vulnerabilities, carries out an assessment and publishes accessible information in a centralised repository. These reports extend back to 2011, when ICS-CERT started to publish advisories. Whilst analysing the dataset of CVEs,

²We validated all these concerns and our motivation for this work in a conversation with an ICS Cyber Security Professional working in a large organisation dealing with relevant information as part of security operations.

we found many issues with the data quality. The largest issue is that of inconsistencies both within CVEs, and across different CVEs. In particular the listed Common Platform Enumerations (CPEs) often do not match the CVE description, and a single device may have multiple different CPEs with variations on how device names are represented. Because this variance could cause issues whilst searching for CVEs for a particular device, we decided to perform an in-depth analysis into these inconsistencies. This leads to two research questions that we aim to answer:

- (1) How long are Siemens vulnerabilities “in the wild” before being discovered and assigned to a CVE?
- (2) How accurate is CVE information to an asset owner and how much does the information within CVEs vary?

Based on these questions, we investigate how asset owners (and vendors) can be better informed for risk and security management in OT environments. This is achieved by carrying out the first significant study on these critical issues, focusing in the first instance on Siemens ICS vulnerabilities, and identifying how any issues identified can be mitigated or resolved. Our contributions are:

- The first significant study of Siemens ICS vulnerabilities,
- Analysis of the time a vulnerability existed ‘in the wild’ before CVE publication,
- A detailed review of CVE accuracy and how this can affect risk management in assets and infrastructure,
- Development of guidance to reduce vulnerability exposure windows and improving the accuracy of ICS vulnerability reports.

This paper is structured as follows: We give an overview of ICS systems, vulnerability management and related work in Section 2. In Section 3 we give an overview of our methodology for performing the analysis, followed by an overview of the data in Section 4. We then discuss the lifespan of vulnerabilities in Section 5. Section 6 presents an analysis of the various issues we found within the CVE dataset, followed by a discussion of potential solutions in Section 7.

2 BACKGROUND AND RELATED WORK

2.1 Background

2.1.1 Industrial Control Systems (ICS). ICS defines the devices, software and services that provide control and information to physical processes in factories, power plants, water treatment facilities and other parts of critical national infrastructure (CNI). They include a wide range of devices such as programmable logic controllers (PLCs), human-machine interfaces (HMIs), remote telemetry units (RTUs) and countless other support devices including those for networking and security appliances. Major manufacturers of such systems include Siemens, Allen-Bradley, ABB, General Electric and Honeywell.

Unlike typical IT systems, ICS devices are traditionally designed for safety and reliability. A device, once installed, should continue operating for a number of years, sometimes decades, and be robust against external factors such as electromagnetic interference. Security, on the other hand, has only become a major consideration for vendors and asset owners since well-publicised attacks like Stuxnet. These devices are intended to be isolated from the internet and wider corporate networks, however, increasingly this separation

is being blurred. A casual search on Shodan reveals hundreds of ICS devices visible to the public internet. Additionally, the advent of Industry 4.0 has introduced internet-of-things (IoT) and cloud components into industrial systems, providing new methods of entry to attackers [6, 15].

It is well-known that industrial devices are often slow to receive patches as raised by Wang et al., measuring patching behaviour of over 100,000 industrial devices publicly visible on the internet (through Shodan), and finding that around 50% patch newly disclosed vulnerabilities within 60 days [23]. This is partly due to a reluctance to take processes offline whilst updating, requirement for safety assurance and the potential for errors to appear in currently working systems. This lack of patching means that it is important that asset owners and operators know which devices within their systems are affected by vulnerabilities, so mitigations other than patching, such as firewalls and intrusion prevention systems, can be utilised. As an example, operating an $n-3$ policy, with a device kept at most 3 firmware version behind the latest, in Siemens S7-1500 products can equal a year in published firmware updates, with $n-2$ being over 6 months, driven by safety requirements and exhaustive testing by the asset owner before deploying the update.

2.1.2 Vulnerability Management. A Common Vulnerability Enumeration ID (CVE-ID) is a unique reference to a vulnerability, issued by a CVE Naming Authority (CNA). The maintainer of the CVE database and primary CNA is the MITRE Corporation. Many vendors operate as CNAs for their own products (e.g. Siemens), or act as a third party for the issuing of CVEs relating to the products of other vendors (e.g. CERT@VDE). A CVE consists of a ID number, such as ‘CVE-2020-1234’, where ‘2020’ is the year the ID was allocated and ‘1234’ a unique number (per year) for the vulnerability. A CNA may reserve multiple CVE IDs in advance, and not all CVE IDs will be used. Due to delays in publication and early reservation of IDs, a CVE published early during the year may have a year component before the actual year of publication. A CVE also usually features a plain-text title, description (which lists affected products and gives a high level description of vulnerability), a list of references (such as links to vendor alerts) and a creation date. Further, through the NIST National Vulnerability Database (NVD), CVEs can be matched to a list Common Platform Enumeration (CPE) IDs, which represent affected products in a “structured” format.

The CPE is a way of describing hardware, applications and operating systems (including firmware) in a machine-readable way, where the NVD CVE record states the list of CPEs, start and end versions of the vulnerability. CPEs are stored in a dictionary maintained by NVD, who attach these to a CVE if a product is affected.

As an example, the CPE Vector `cpe:2.3:siemens:simatic_s7-1500_firmware:*:*:*:*:*` can be broken down as follows – CPE Version 2.3 (issued in 2011), the ‘Part’ is an operating system (:o:), the Vendor is Siemens, the Product is the SIMATIC S7-1500 firmware and any Version, Update, Language, Edition, Language, Software Edition, Target Software, Target Hardware is affected.

Vulnerabilities are also assigned a Common Vulnerability Scoring System (CVSS) score, from 0 to 10. The CVSS score can be summarised as “Low” (0.1-2.9), “Medium” (4.0-6.9), “High” (7.0-8.9) and “Critical” (9.0-10.0). The purpose of the score is to both assign the

vulnerability a severity, as well as provide a measure for prioritising responses to vulnerabilities. These scores are calculated by NVD. As part of the CVSS vector, the impact on confidentiality, integrity and availability is also provided, rated as “None”, “Low” or “High”.

2.1.3 EU Network and Information Systems (NIS) Directive. The NIS Directive aims to improve the cyber security of infrastructure systems [22]. This placed a responsibility on the asset owner to appreciate their security posture and drive an improved security culture when operating essential services. In the United Kingdom, the NIS Directive implementation is guided by the Cyber Assessment Framework (CAF) [13], which has indicators of good practice, by which an asset owner is measured on whether they have not achieved, partially achieved, or fully achieved the objective. Objective A3 focuses on effective asset management, where A3.a requires up-to-date assets in order to be fully compliant, and B1.a encouraging a continuously improving security and risk management regime.

2.2 Related Work

Security incidents and vulnerability reports have been previously assessed at a high level, specifically in the area of the types of vulnerability that exist in ICS, with no longitudinal surveys of ICS CVEs to date. The last ICS-CERT report in 2016 [10] does not review how long ICS owners could have been exposed to vulnerabilities. Hemsley and Fisher provide a timeline of high-impact ICS incidents, but do not assess the ‘life’ of the vulnerability which allowed these significant events to take place [9]. Thomas and Chothia assess the types of vulnerabilities persisting in ICS and their detection, but do not include temporal analysis or assess the correctness of CVE information as they use ICS-CERT as a ‘single source of truth’ [20]. In IT solutions, where source code is available, analyses against Google Chrome, Drupal, PHP-MyAdmin and Moodle have been used as case studies [14, 21], where in Chrome [14], the distribution of time between release and vulnerability identification is assessed where most vulnerabilities were ‘in the wild’ for between 10 and 40 months. This is very different in Walden et al.’s survey [21], where the distribution time was approximately 4000 days (131 months), with a median of 831 days (27 months). This was possible because the applications under assessment were open-source, and thus, given a vulnerability, the commit that introduced that vulnerability could be identified. In ICS, however, this is generally not the case, as these products exist as proprietary, closed-source systems. In 2016, Kaspersky Lab’s analysis of ICS vulnerabilities [1], found that 85% of vulnerabilities were patched or resolved, but 5% were only partly resolved, and 6% were not patched as the affected component was either no longer sold, or supported. In a new threat landscape, published in 2020 [12], some analysis into how long a vulnerability existed was conducted, but kept at a high level, showing how a vulnerability propagated through common, shared, components. An assessment by Dragos of the correctness of CVSS scores in ICS vulnerabilities found that in 2018, 32% contained errors [4], improving to 19% in 2019 [5]. This demonstrates how asset owners may not fully appreciate the criticality of an incident, however, no comment is made on the accuracy of affected devices in their analysis. Incurring delays in the assessment of CVEs also delays the ability of an asset owner to understand the severity of the




Figure 1: ICS vulnerability data import workflow.

vulnerability. A study by Ruohonen in 2019 shows that the average time to assess and assign a CVSS score to a CVE is 134.7 days [16]. Whilst ICS-CERT and vendors may include this score, this will not be included in the CVE until assessment by NVD.

3 DATASET AND METHODOLOGY

Data Sources and Building the Dataset: ICS vulnerabilities are published by a number of authoritative sources, for example in vendor advisories, CVEs and ICS-CERT alerts/advisories. Given the varying detail, formats and accessibility of vendor advisories, using vendor advisories alone is not suitable for such analysis.

Our ICS vulnerability dataset³ is curated from 1,143 ICS advisories, 2,327 CVE entries and 44,030 CPE listings using a similar process proposed in [20] to unify ICS vulnerability information. This dataset supports our analysis of ICS vulnerability reports, where previous work has reviewed the accuracy of vulnerability reports, in particular for the stated CVSS scores and vectors [4, 5]. It is built using data from a number of sources, primarily ICS-CERT Advisories, MITRE, the National Vulnerability Database (NVD) and vendor websites. We use ICS-CERT as our main source of information, where we follow references to CVEs and CPEs to extract further information, using the vendor website to extract key dates, and harmonise CPEs to their respective product portfolio range. Figure 1 shows how this data is collected, imported and curated. For each source, a brief overview and import workflow is as follows:

ICS-CERT Advisories: These are published by the US Cybersecurity and Infrastructure Security Agency (CISA), responsible for providing authoritative vulnerability information to ICS asset owners and the wider community. To the best of our knowledge, this is the most comprehensive and complete source of ICS vulnerability information, where we can focus on ICS-specific vulnerabilities, which other services (e.g. CVEDetails) do not highlight. ICS-CERT advisories are published in HTML format, which we convert into markdown to ease parsing and aid the extraction of CVE numbers.

MITRE CVEs: The MITRE corporation is responsible for the CVE scheme, where CVE numbers are allocated to CNA and represent a single point of numbering. MITRE CVEs are available as individual JSON files for parsing, where some CVE numbers may never be assigned. From MITRE, we extract the CNA assigner, the organisation that requested the CVE number and contributed the details to the CVE record.

NVD CVEs: The NVD, provided by NIST, appraises and analyses vulnerability reports (where MITRE is not considered authoritative

for vulnerability information), defining a CVSS score and CPE list. NIST CVEs are provided in regularly-updated year-based JSON files, containing all assigned CVEs, which we split to simplify the lookup process rather than parsing an entire year’s CVEs to get a single CVE. When a CVE ID has been identified in an ICS advisory, the corresponding MITRE and NVD CVE files are retrieved and the Common Product Enumeration (CPE) list, CNA, and vulnerability information are imported.

Vendor Websites: As our study also considers how long vulnerabilities have existed since the initial release, we use published release notes and dates stated for the affected versions from the vendor website. This, however, can be a limiting factor, where some vendors require a support contract to access this information or archive such information after a period of time. This information is manually entered into the dataset as it requires careful validation of the presented information. For the purposes of this paper, we used the Siemens Industry Mall and Support websites to obtain product release and firmware update dates, which were manually entered. These sources were also used to unify similar CPEs to a product range where a CPE may have stated a specific model within a product range, and the product range itself was also affected. This enables us to consider the most affected product ranges.

Of the 2,327 CVEs in our ICS vulnerability dataset, 424 had “Siemens” as the associated vendor. We removed any CVEs which only feature application CPEs. This is because we want to focus on ICS devices, with CPEs representing hardware and firmware more likely to cover these CVEs. This was achieved by removing CVEs in which the CPE list only contains “a” as the part field, and retaining CPEs with the operating system (“o”) or hardware (“h”) type. This resulted in a final device dataset of 207 CVEs for analysis.

3.1 Temporal Analysis

In order to state how long a vulnerability has existed, we require reliable information from authoritative sources, so as to not skew our results.

For our temporal analysis, whilst a majority of data used for our analysis could be collected, parsed and stored through automated means, there was some manual effort required to supplement the dataset, in particular, adding release dates of affected software and data validation. ICS-CERT and the CVE records often include references to the vendor advisory and state appropriate mitigation/patches, but do not state when that was made generally available. Therefore, manual review of the vendor website was undertaken to establish when the affected versions and updates were published. Due to the significant level of variation on the websites (e.g. date formats and location of information), manual effort is

³The raw, unprocessed, CVE dataset is available at <https://github.com/uob-ritics/cpsiotsec2020-dataset>.

more effective and reliable than automatic extraction, especially for date formatting. Automated efforts also lack expert contextual knowledge when identifying the initial release, e.g. when new generations of a product are introduced, but have a common name.

For each of the CVEs marked for analysis, we reviewed which versions were affected in the CVE description. If the CVE stated that all versions of that product were affected, we searched for the sales release date, as that vulnerability would have existed when the product went to market. If a start version was stated, we attempted to find (typically through product notes) when that particular firmware was issued. For a very small number of CVEs, we were not able to establish the date of release for the stated start version. In these cases, we looked for the earliest version issued within the range of affected versions (i.e. if a CVE stated that versions 3.0 through to 3.5 were affected and we could not establish when 3.0 was issued, but a date for 3.1 was available), we use the date of that version to indicate that the vulnerability existed since at least that date.

We made the decision to ignore CVEs related to BACnet products, largely covering building control systems, as these are not presented on the main Siemens Industry website, and do not have sufficient temporal data available. We also removed CVEs affecting IP cameras originally produced by Siemens (Vanderbilt acquired this portfolio in 2015), and CVEs where the descriptions and/or CPEs were too vague to identify specific products. This accounted for 36 CVEs within our dataset.

From this manual review process, we were able to collect data for 165 CVEs (the “temporal” dataset). For some CVEs, we were unable to verify dates for all CPEs, however as long as there is at least one date available we include the CVE in our analysis.

Issues: When attempting to backfill key dates using the vendor website, we noted significant variances in how support documentation was written. As an example, in one product line of PLCs, Siemens had more than one document containing the full list of firmware releases, where one contained a fuller firmware history than the other. Notably, this can also vary by which localised version of the website is used. For another popular range of PLCs by Siemens, significantly less information was available, with fewer versions of firmware displayed, and no release dates held in a single document, rather fragmented over a number of download pages.

As stated earlier, it was not possible to obtain definitive dates for when affected versions of firmware was issued for 36 CVEs. The products identified in these CVEs were typically older, outdated or no longer supported, with release notes unobtainable from authoritative sources. For some products, Siemens did not state the product lifecycle (sales release, delivery release, discontinuation), or the product was no longer part of Siemens’ portfolio, now owned by a subsidiary, where support contracts were required. Where this occurred, we attempted to find evidence of the first available firmware version, which would indicate that the product existed at firmware publication. In other cases, we do not state a ‘start’ date.

We also noted the formatting of release dates varied both within articles and across similar documents for the same product (e.g. DD/MM/YYYY, MM/DD/YYYY and dd.mm.yyyy) where multiple articles had to be consulted to ensure the date range was valid for the format we use (DD/MM/YYYY).

3.2 Inconsistency Analysis

In validating the correctness of CVE information, we assessed the full set of 207 Siemens CVEs representing hardware and firmware vulnerabilities, using the CVE description and CPE list to validate whether the CVE description matched the list of CPEs and vice-versa. We refer to this as our ‘inconsistency’ dataset.

As part of this assessment, we extracted the list of products named in the CVE description, the individual CPEs for that CVE and verified whether there were any missing products in the list of CPEs, any products missing in the CVE description (where there was a CPE entry) and also that the firmware versions stated were contained in the CPE (or were at least valid). Manual work to unify CPE details was undertaken to harmonise formatting, and identify the specific models and versions affected. As part of this exercise, we carried out taint analysis to determine the scope of the CPE vectors, and the variations that occurred between CVEs affecting the same product.

We also observed that between the NVD and MITRE CVE feeds, the CNA assigner varied between the two, where NVD would attribute all CVEs as having been assigned by the MITRE Corporation. As a result, we use the MITRE feed to clarify who assigned the CVE number, where some statistics are given in Section 4.

To determine how accurate the CVE descriptions were, we carried out a manual analysis of every CVE in the dataset. Each CVE was listed alongside every CPE associated with it. The number of devices listed in the description was counted and compared to the number of unique CPEs. Where more than one CPE existed for the same device (e.g. with firmware versions), they were grouped together and counted as 1 device, as we only care about devices missing from the CPE list. Additionally, if a firmware version was stated in the CPE, we verified if it was in the CVE description.

4 OVERVIEW OF DATA

In this section, we provide an overview of the data held in our dataset which we will use to analyse how long Siemens vulnerabilities existed before CVEs were issued, and how accurate the vulnerability reports were from the perspective of the CPE vector. This provides additional evidence to previous work [4, 5] that the accuracy of ICS CVEs is not necessarily always correct.

As shown in Table 2, there are a total of 424 CVEs relating to Siemens industrial products. The majority, 288 out of 424 (67.9%), have been assigned since 2016.

4.1 Issuing CNAs

Table 2 shows the number of CVEs issued describing Siemens industrial products per year, divided into into the number issued by each CNA. We observe that Siemens only started issuing CVEs themselves in 2016, issuing CVEs for their own products (i.e. became a CNA) with the MITRE Corporation largely responsible before then. As Siemens is its own naming authority, a vulnerability needs to be reported to Siemens, as MITRE will only issue a CVE number if the vendor does not have an assigned CNA.

As we observe from Table 2, although Siemens is the predominant CNA for vulnerabilities within their portfolio, there are CVEs raised by other CNAs, such as Intel, HPE, and RedHat, where vulnerabilities in shared components and libraries have been found and

	2010	2012	2013	2014	2015	2016	2017	2018	2019	2020 (Jan - Apr)	Total
cert@cert.org	0	19	1	0	0	0	1	3	0	0	24
cve@mitre.org	1	1	16	33	34	30	1	1	9	0	126
ics-cert@hq.dhs.gov	0	15	16	0	0	0	4	0	1	0	36
productcert@siemens.com	0	0	0	0	0	7	36	50	121	16	230
secalert@redhat.com	0	0	0	0	0	0	0	0	2	0	2
secure@intel.com	0	0	0	0	0	0	0	3	0	0	3
security-alert@hpe.com	0	0	0	0	0	0	0	0	2	0	2
talos-cna@cisco.com	0	0	0	0	0	0	0	0	1	0	1
Total	1	35	33	33	34	37	42	57	136	16	424

Table 2: Number of CVEs per year relating to Siemens industrial product by issuing CNAs

Siemens products were identified as vulnerable. The one exception to this, however, is the MITRE Corporation, which issued 9 CVEs linked to Siemens products in 2019. Upon review of these 9 CVEs, they were issued after a number of vulnerabilities were found in Wind River’s VxWorks real-time operating system, affecting F5, Sonicwall, Netapp and Siemens. In 2019, Intel also issued 3 CVEs (affecting Intel AMT), which included Siemens industrial PCs.

Finally, we have the case where Cisco Talos issued a CVE in 2019, CVE-2018-3991, identifying a vulnerability in WinCC Open Architecture, which, whilst Siemens-branded, is operated by a third-party, where Siemens does not have to be the named CNA and MITRE can therefore issue the CVE.

4.2 Number of Vulnerabilities

As seen in Table 2, the number of CVEs issues for Siemens products was consistently between 33 to 35 from 2012 to 2016, and has since steadily increased year on year, with a particular noticeable jump from 57 to 136 in 2019. As of April 2020, there were only 16 CVEs issued for the year. This low number could be due to the time it takes to report, triage and publish vulnerability reports, as well as the COVID-19 pandemic, which has impacted productivity across a number of sectors, including research facilities, resulting in fewer reported vulnerabilities.

4.3 Number of Devices per Vulnerability

For the full set of Siemens CVEs and all CPEs, including those only affecting applications, there is a mean of 8.27 CPEs (SD 16.09) listed per vulnerability, with a maximum of 148. If we focus on the 207 CVEs in our final dataset, with only hardware and firmware CPEs, we have a mean of 11.97 CPEs (SD 20.87), with a maximum of 148. Of those CVEs, only 14 apply to a single CPE. The largest vulnerabilities, CVE-2017-2671 and CVE-2017-2681, both cover 148 CPEs over a wide range of devices, as they both relate to vulnerabilities in PROFINET, a protocol common to many Siemens industrial devices. It important to note that just because a CVE has a large number of CPEs, it does not necessarily mean that a wide range of devices are affected, as many vulnerabilities affect multiple versions of the same base product, which can each be assigned a CPE, though, in general, the greater the count of CPEs, the more product ranges are affected by a vulnerability.

4.4 CVE Severity

Of all 424 Siemens CVEs in our master dataset, the mean CVSS score is 7.07. In our reduced dataset of hardware and firmware only CVEs, the mean CVSS score is only slightly higher at 7.25. Both of these indicate a “High” impact on average. As shown in 3(a), the full

dataset contains 57 “Critical” vulnerabilities (a CVSS score >9.0), or 13.4%, whilst in the reduced dataset there are 22, or 10.5%.

Table 3(b) shows the impact of the CVE on confidentiality, integrity and availability. In ICS systems, availability is generally the most important impact to consider (behind safety), followed by integrity and then confidentiality. As seen in the table 3(b), in both datasets availability has the greatest volume of “High” impact CVEs, followed by confidentiality and then integrity. This is somewhat expected, as availability (denial-of-service) and confidentiality (information leakage) vulnerabilities are generally easier to identify.

CVSS Severity	All Siemens	Siemens Hardware and Firmware
Critical	57	22
High	194	110
Medium	17	72
Low	156	3
Total	424	207

(a) CVSS Severity

	All Siemens			Siemens Hardware and Firmware		
Impact	I	A	C	I	A	C
High	154	235	171	66	125	73
Low	80	44	109	40	18	45
None	190	145	144	101	64	89
Total	424			207		

(b) CVE Impact on Integrity (I), Availability (A) and Confidentiality (C)

Table 3: CVE severity and impact

5 HOW LONG DO VULNERABILITIES EXIST BEFORE A CVE IS GENERATED?

Understanding the time in which an asset owner may have been exposed to a vulnerability enables effective risk management and response. In this section, we analyse the time of a vulnerability entering the product ecosystem to the CVE being published.

5.1 Results

As discussed in Section 3.1, we were able to identify temporal data for 165 of our 207 device-related CVEs.

Table 4 presents the overall results of our analysis. We find that, on average, vulnerabilities existed “in the wild” for 1897 days (5.2 years), with the maximum time between release and CVE publication being 5352 days (14.7 years) and the minimum, 115 days. In the worst case, a vulnerability was in the wild for up to 8152 days (CVE-2019-10936 and CVE-2019-10923). This timeframe, however, must account for any responsible disclosure, triage, investigation

and coordination, with the time to discover a vulnerability potentially being lower. When a vulnerability is reported, there is no guarantee the latest version was reported, where the vendor may carry out a detailed review of all newer revisions at the same time. Our analysis, however, uses what is stated in the CVE – if a more recent version was identified as vulnerable, it only confirms how long a specific vulnerability persisted before a CVE was published. Figure 2 shows the full timelines for 139 of the 165 CVEs group by the stated Common Weakness Enumeration (CWE).

Part of the reason for these long lifetimes is that on disclosing a vulnerability, the vendor will check all previous versions of the firmware for the vulnerability. We find that in the majority of cases, vulnerabilities exist in all versions of the firmware before the current (or a recent) version. For many devices which are generational, whilst later generation devices are only capable of running newer firmware versions, past generations, whilst no longer available to purchase though running older firmware versions may still be in operation in large volumes, which are also vulnerable. As we have shown, vulnerability disclosure has increased in recent years, with more vulnerabilities spanning generations being discovered. Moreover, the longer a vulnerability persists, the installed consumer-base increases, with a higher number of potentially exposed devices requiring swift patching when a CVE is issued. We do not test whether the patches resolved the vulnerability, a factor which could affect the lifetime of vulnerabilities, where we must trust what is stated in the vendor advisories as the resolution.

5.2 Relationship between Vulnerability Lifetime and Type

As part of our analysis, we wanted to see if there is any relationship between the type of vulnerability, (the CWE assigned to the CVE), and the lifetime of the vulnerability. For this purpose, we use the CWE as defined in the ICS-CERT advisories as these are usually more precise to the actual vulnerability, as they are produced with greater input from vendors. The results of this mapping are presented in figure 2. Within the dataset, there are a total of 61 CWEs. We focus only on CWEs that have at least 2 associated CVEs, leaving 29 CWEs across 139 CVEs, with an overall mean lifespan of 1921 days (5.26 years).

The most common CWE is CWE-20 (improper input validation), which is not unexpected due to improper input validation being a root cause of many other vulnerabilities, such as injection attacks and buffer overflows. These CVEs, on average, take 2111 days (5.8 years) to be discovered, slightly above average.

We observe that there is a large difference between the shortest lifetime vulnerability (CWE-710, 205 days) to the longest lifetime vulnerability (CWE-319, 4294 days). The two CVEs relating to CWE-319, CVE-2018-13808 and CVE-2019-10926, both affect different product ranges and were discovered in different years.

From Figure 2, similar types of vulnerabilities vary greatly in vulnerability lifetime. CWE-80, for example, covering basic cross-site-scripting (XSS) vulnerabilities, has an average lifetime of 1125 days, whilst CWE-79, also covering XSS vulnerabilities has an average time of 2611 days.

Of particular interest are examples of vulnerabilities that we would expect to be easier to discover are some of the slowest to be discovered. For example, CWE-264, CWE-284 and CVE-306 all

	Mean	Std Dev	Min Value	Max Value
Worst Case	2421	1717.9	115	8152
Best Case	1398	1095.2	5	5352
Average	1897	1178.6	115	5352

Table 4: Mean time taken (in days) between a product going to market/firmware being released and the CVE being issued, in the best case, where we take the closest date to the CVE release, the worst case, where we take the most distant date per CVE, and the average case, which is the mean of all dates found for a CVE.

relate to access control vulnerabilities, but these take more than 2000 days on average to be discovered. CWE-200, CWE-522 and CWE-319 all relate to exposed sensitive data, with average lifetimes of 1982, 2037 and 4294 days respectively, or at least 5.4 years.

5.3 Relationship between Vulnerability Lifetime and Severity

We wanted to measure if there is any relationship between how long a vulnerability existed and its severity according to the assigned CVSS score. For the 165 vulnerabilities for which we were able to identify dates, the average CVSS score was 7.23, with four being assigned the highest score of 10. To measure the correlation, we calculated the Pearson correlation coefficient, where ± 1 indicates strong correlation, between the average vulnerability lifetime and CVSS score. The Pearson correlation was only -0.024, indicating little to no correlation between vulnerability age and severity.

Further, we found that there was little relation between the age of the vulnerability and the assigned impact on integrity, availability or confidentiality.

6 INCONSISTENCIES IN VULNERABILITIES

Knowing how long an asset owner has been potentially exposed to a vulnerability is valuable for incident response and managing assets. Given the range of automated tooling for security-relevant information to be presented to an asset owner, e.g. OpenCTI⁴, which allow an asset owner to monitor CVE and open source intelligence reports for assets they own, the quality of the data being produced is equally important. If this information is incorrect at the source, its propagative effect can lead to an uninformed asset owner.

6.1 Difference between CVE Description and Listed CPEs

During data-cleansing and validating the quality of data, we cross-referenced the CPEs assigned to a CVE with the CVE description, and any product security advisories published by the vendor. Given the shift to automated alerting of potential vulnerabilities by ‘watching’ a supplied list of assets, there is a risk that the CPE list could not contain all affected devices, resulting in no alerts to an asset owner. Alternatively, if the CPE list is too open, i.e. a product range and no specific model number/identifier is defined, an asset owner will receive too many alerts which may then be ignored, leading to critical vulnerabilities being overlooked.

From our analysis of 207 CVEs affecting Siemens products, 34 (15.5%) CVEs understated the affected products, where the devices stated in the CVE description (the human readable text) were not

⁴<http://opencti.io/>




Figure 2: Timegraph of ICS CVEs, where a CWE (root cause) had at least 2 CVEs and the range in which the vulnerability could have been ‘in the wild’. + denotes the worst case, × the ‘best case’ and • when the CVE was published. The colour of the line indicates CVSS severity, where yellow is ‘Low’, amber ‘Medium’ and Red ‘High/Critical’.

fully contained in the list of CPEs. Unless the ICS owner searched all CVE descriptions for matching assets, they would not be informed of a vulnerability, as the CPE list would be incomplete. In the worst case, CVE-2019-6568, a vulnerability affecting the device webserver, lists 64 devices in the description, however only lists 46 CPEs – 18 fewer devices. Another example of note is CVE-2019-18336, a severe uncontrolled resource consumption vulnerability affecting the SIMATIC S7-300 CPU family, including the related ET200 CPUs and SIPLUS variants, plus SINUMERIK 840D SL devices and SIMATIC TDC CPU555 devices. In this vulnerability, a specially crafted packet will cause the device to enter a failure state, requiring a restart. The CPE list, however, only covers multiple specific S7-300 variants and the SINUMERIK 840D SL device, meaning that the popular ET200 range and SIMATIC TDC CPU555 are not covered. 30 CVEs (14.5%) featured greater precision in the CPE list than listed in the CVE description. Generally, this is due to individual variants of a device being listed in the CPE (such as 1211c, 1212c, 1215c), whereas the description only includes a product family (S7-1200).

Overall, of the 207 CVEs listed, only 143 (69%) featured a CPE list that matched the CVE description. For CVEs affecting ‘all versions’, the CPE would apply a wildcard to the version component (stating

all versions were affected), but, as time passes, this vulnerability may have been patched in a newer version of the software, with the end version not being subsequently updated. For an asset owner who may have just procured new assets, this would highlight invalid CVEs requiring action, where the CVEs in question have been mitigated or resolved in the time between CVE issue and purchase.

Similarly, we also observed inconsistencies where configurations of a product were vulnerable and were stated as such in the CVE description, but the CPE vector would only consider the base product. As a result, unnecessary alerts would be raised, where the systems in an asset owner’s environment may not be vulnerable, as they do not have the specified option which results in the vulnerability. Additionally, we observed that in February 2020, Siemens carried out a significant update of its advisories, when its SIPLUS/SIPLUS NET range had been identified to be vulnerable to issues also affecting its SIMATIC S7 range of products. The Siemens ProductCERT advisories had been updated, but the CVE and ICS advisory were not updated, creating a disparity between information sources.

6.2 Inconsistency in Device Descriptions

If the quality of vulnerability information at source is poor, there is a propagative effect that reduces the effectiveness of that information to the asset owner. Whilst understating the vulnerable devices is an issue, mitigated in part though manual review of ICS Advisories, asset owners need the right information at the right time to enable prompt action.

Beyond Siemens, we observed a lack of consistency in the ways the same product and vendor were represented as a CPE vector. In some cases, the acronym of the vendor was used, e.g. ‘GE’ in place of ‘General Electric’. In another one case, as well as the vendor “Honeywell”, we noticed that 2 CVEs refer to “Honeywell” (CVE-2019-18226 and CVE-2019-18228). In both of these CVEs, the same erroneous CPE can be found (cpe:2.3:o:honeywell:h4w2per2_firmware:-:...). and both CVEs were issued on the same date (31-10-2019). It is likely that this is due to human error, with two related vulnerabilities published simultaneously with an identical list of CPEs. In the former case, an asset owner can monitor for alternative representations, whereas the latter would never be detected, and lies solely on quality control when a CVE is prepared for release. For the two vendors, Schneider Electric and Phoenix Contact, there were 4 different representations each (including typographic errors). This demonstrates that these issues exist across vendors.

The scale of variation in vendor representation was minimal compared to product representation. As an example, the Siemens ET200S PLC had two different representations, et200s and et_200s. If following the Siemens Industry Mall designation, 4 CVEs would have been missed, as they were represented as et_200s, whereas 31 followed the correct designation. In the case of the S7-1200 range, s7-1200_cpu, simatic_s7-1200_cpu and simatic_s7_cpu_1200 have been used, in addition to the individual product names (e.g. 1212C). We also observed that, for example, in CVE-2019-19278, the CPE included MLFB, a term to describe the Siemens product ‘ordering’ code. In the corresponding Siemens Security Advisory and CVE description, it states this as the product code prefix for products affected. In the same CVE, the description was explicit that a specific option had to be fitted to the affected system in order to be vulnerable. The CPE, however, does not include this, therefore suggesting all products were affected, when, in reality, only a subset was.

Where Siemens was not the CNA for a CVE, e.g. CVE-2019-12259, the CVE description and list of CPEs does not match what Siemens ProductCERT states is vulnerable. In the CVE, only the SIPROTEC 5 is stated, when the Siemens Security Advisories identify additional products affected (RUGGEDCOM and two series of power meters). In the case of this CVE, NVD last carried out a review in September 2019, where it is awaiting reassessment, where Siemens included the affected RUGGEDCOM products when the CVE was published, and the power meters were added in June 2020.

When variations of how products are represented, there is a risk that an operator may miss critical vulnerabilities due to a mismatch in the expected CPE. We noted in some CVEs, the core model number/product identifier was used, identifying the exact device affected, where we explore its utility further in Section 7.

CVE ID	Unique CPE Count	CVE Updates	CVE Description Changed?	CPE List Updated?	SSA Revised?
CVE-2019-13927	16	0	×	×	×
CVE-2019-13940	23	3	✓	×	✓
CVE-2019-6571	24	1	×	×	✓
CVE-2019-6584	24	1	×	×	✓
CVE-2017-12741	38	15	✓	×	✓
CVE-2019-10923	46	4	✓	×	✓
CVE-2019-6568	46	11	✓	×	✓
CVE-2019-13946	51	2	✓	×	✓
CVE-2019-10936	54	4	✓	×	✓
CVE-2017-2680	74	14	✓*	×	✓
CVE-2017-2681	74	9	✓ ^o	×	✓
		64	8	11	10

* - The description was modified, removing all affected devices, and replaced with a generic statement with excluded devices.

^o - The CVE description was modified to be generic, similar to CVE-2017-2681.

Table 5: CVEs which affected more than 15 unique Siemens products, showing the number of times the NVD CVE record was updated since publication, if the CVE description or CPE list were updated, and whether the Siemens Security Advisory (SSA) had been revised since the CVE was published.

6.3 Reasons for Inconsistencies

As our study focuses on CVEs affecting Siemens’ industrial automation portfolio, we must consider potential reasons for the variations in CPE vectors.

In addition to carrying out formal assessments of CVEs, NVD operates the ‘CPE dictionary’⁵. This contains CPEs, both from existing CVEs as well as contributed CPEs from the open community. Here, references for a given CPE can be supplied, such as source information of how that CPE was constructed (e.g. from a product page). In the case of the Siemens CPEs we sampled in the Dictionary, all references are to Siemens ProductCERT advisories.

When a CVE is published, the NIST NVD assessment takes place - only the CVSS and CWE information from the naming authority may be used⁶, and no CNA was named as being a contributor to CPE information. Instead, NVD staff will populate the affected products (CPEs) for a CVE. Whilst NVD is formed of experts, this is the source of variation in CPE vectors, leading to some of the observations we give in this section, where similar products were represented very differently and increasing the opportunity of an asset owner being misinformed. This, however, makes this an industry-wide issue for all CVEs, where the vendors may be in a better position to propose CPEs, as this vital context ensures that all information is accurate, consistent and well-represented.

Table 5 illustrates a few example cases where the CVE description or Siemens advisory was updated, but the CPE list was not subsequently updated. In each of these cases, the CVE description had been updated, flagging that the CVE required re-assessment by NVD, but even in the case of the oldest CVE in the table, CVE-2017-2680, the only assessment was in May 2017, where the CVE has had 9 significant changes to its description since then, but the corresponding CPE list has not been updated. This means that when new products are introduced in scope, the asset owner, watching the CPE list, would never be informed, as the CPE list was not updated.

⁵<https://nvd.nist.gov/products/cpe>

⁶The process undertaken is given at <https://nvd.nist.gov/vuln/cvmap>.

6.4 Effect of Inconsistencies

In order to show the potential impact of these inconsistencies, we take a number of devices that we own, and perform a search on different terms to identify the variations in the way assets we own were represented. We name a few products we own and the product component of the CPE vector that was assigned:

- Siemens LOGO! PLC: 6ed1052-1md08-0ba0, logo\!8_bm_firmware, logo\8_bm_fs-05
- Siemens SIMATIC S7-1212C/S7-1211C: simatic_s7-1200, simatic_s7-1200_cpu_1212c, s7-1200_cpu_1212c, simatic_s7_1200_cpu, 6es7211-1ae40-0xb0
- Siemens SIMATIC S7-1518-4 PLC: simatic_s7-1500_cpu_1512c, simatic_s7-1500_cpu, simatic_s7-1500
- Siemens SIMATIC ET200s: simatic_et_200s, simatic_et200s

From these examples from assets we own, the level of variation impacts our ability to verify whether we have correctly defined all the possible CPEs to match against, in addition to having sufficient detail to determine if our assets are vulnerable. One particular example of interest is how the Siemens LOGO! range has been represented, Siemens' small-scale automation portfolio. Where the product code was given, we could easily determine if we did have the specific PLC, whereas the two CPEs containing logo\!8_bm were too vague, as BM is 'base module', where a number of generations of the LOGO! range have been developed and sold, some vulnerable and others not due to firmware and design changes, however retaining the same product name, e.g. LOGO 12/14 RCE.

As identified earlier in this section, some CPEs assigned to a CVE overstated the scope of affected systems. With CVE-2019-6578, if the product range was stated, all asset owners would be alerted to the vulnerability, when a specific configuration option had to be selected in order to be vulnerable. If an asset owner were to monitor for their product, they could receive a false-positive alert, where their system is not vulnerable. It, however, has the benefit of allowing them to apply discretion to determine whether they were affected. At the same time, it would lead to a higher volume of alerts which could lead to critical vulnerabilities affecting that asset owner being overlooked due to the amount of 'noise' produced, where targeted alerts are not being given. In the case of vendors incorrectly being defined, an asset owner would never be informed of applicable CVEs as the vendor in the CPE would not match what the asset owner would be looking for.

Further to the results shown in Table 5, during the time of this paper being drafted and the dataset being validated, we observed changes to the CVE description. Previously, the CVE description contained the list of affected devices and versions, enabling an asset owner to carry out textual search for vulnerabilities. As an example, CVE-2017-2680 and CVE-2017-2681 were updated in July 2020 with a much shorter CVE description given, presenting a high level overview of the vulnerability. At the same time, CVE-2019-10943 provided contradictory information to asset owners in the CVE description, specifically whether the SIMATIC S7-1518-4 was affected. In the CVE description, it first states only the MFP variant is excluded, but immediately after it states both the PN/DP and MFP variants are excluded (i.e. all S7-1518-4 products as they are one of PN/DP, MFP or both). This lack of clarity can convince an asset owner that they do not need to take any action, as the

latter statement would be considered the most recent statement. On the Siemens Advisory, however, the PN/DP variant exclusion was removed in March 2020, where this confusion could have serious consequences. The vulnerabilities in this advisory allow the integrity protection to be circumvented, which could impact safe operations.

7 IMPROVING ICS VULNERABILITY INFORMATION

From our analysis of ICS CVEs, taking Siemens as a case study, we have established that there is some considerable time between devices and firmware being released to a vulnerability being discovered, reported and published. It is important to note that the time difference will include the reporting, triage and assessment time. Moreover, we find that the consistency of CVE descriptions and statements of affected products is not where it needs to be in order to support asset owners of critical national infrastructure in managing risks in an informed manner.

In this section, we outline the potential guidelines which can be implemented at various levels, including vendors, NVD and NIS Competent Authorities. These will not only improve the security of ICS systems through improved testing regimes, but also the quality of data available to ICS owners to make appropriate security decisions at the right time with the right information.

7.1 Reducing Vulnerability Lifespan

7.1.1 Better Defined Testing Strategies and Certifications. It is well understood that industrial control devices undergo rigorous testing to ensure the devices will operate continuously and safely, for an extended period of time. Devices require thorough testing and certification before they can be sold. This is not the case, however, for the security of the products. We argue that vendors should employ similar levels of rigour when it comes to the cyber-security of their devices. This can be achieved through the development of industry-accepted testing strategies for devices, including penetration testing of devices before release, as well as the development of certifications that verify devices have been tested to a sufficient degree before release. As we show in Section 5.2, there are multiple vulnerabilities types, such as cleartext private information and improper access control, that have extended lifetimes however should be easier to discover with more thorough testing. Of course it is inevitable that some vulnerabilities will still exist past the vendor testing stages, however vendors can encourage and support third party researchers to identify vulnerabilities to assist in catching vulnerabilities not caught by vendor testing (see Section 7.1.2).

One example of certification already in use is the Achilles System Certification (ASC) program, by General Electric (GE). Achilles illustrates the compliance of vendor control system products with cyber-security requirements specified by the IEC 62443-3-3 standard. Achilles assigned a level from 1 to 4, with 1 representing resilience against casual or coincidental violations of security, and Level 4 providing resilience against intentional violations using sophisticated means with extended resources [8]. According to GE, Siemens currently has 408 devices that are Level 2 certified (which covers protection against intentional violation using simple means with low resources, generic skills, and low motivation), however

uptake is lower amongst other vendors [7]. At the time of writing, no vendor has a product certified above Level 2.

7.1.2 Better Motivations and Resources for Vulnerability Disclosure. Whilst there has been an increased focus on OT security issues over the past few years, very few vendors offer financial motivation for reporting security issues. The bug bounty concept has increased in popularity during recent years, and so offering financial rewards for reporting vulnerabilities could motivate greater numbers of people to focus on OT devices and software, in particular with the potential value of exploits on the black market. Further, acquiring ICS devices can be costly, requiring a large initial investment from researchers. A bug bounty can help researchers recover this cost. Vendors may also provide free or heavily discounted devices and software to recognised researchers, facilitating further device analysis. Companies such as Apple have already started programs, providing modified devices to researchers for specific use in finding vulnerabilities, provided the researcher has a proven track record in disclosing vulnerabilities in Apple or similar products [2].

7.2 Enabling Access to Firmware Histories

Unlike traditional IT, where software updates are issued on a regular basis, OT systems have a slower frequency of updates being issued, partly due to the testing and validation that is required for certification and compliance. In the ICS sector, it is a common practice that once a system has been deployed, it will often not be updated [17, 18, 23]. As previously highlighted, operating within n versions of the latest version of software may be effective in the corporate environment, where such revision cycles may be within months of the most recent version. In OT, however, as we observed with some firmware versions, running even within 3 versions of the latest firmware could be equal to over a year (e.g. S7-1500). Without reference dates, an asset owner cannot adequately implement good configuration management practices.

One of the main challenges in establishing when affected firmware was published, in particular when a vulnerability was introduced at some version and later resolved, was that this information was not readily available, or sparsely distributed, an issue also identified in [4]. For an asset owner, the lack of visibility of firmware history leads to a false view and confusion that the firmware that is running in the live environment is recent. This is particularly critical for those devices which are deemed ‘legacy’ as this enables the asset owner to consider appropriate mitigations that are required if that device is essential in a process.

Access to this history should be possible to recover and publish in a consistent format such that an ICS owner can search for their asset and the firmware history in a ‘change log’ format. Through this format, key changes can also be highlighted, such as critical patches to vulnerabilities, or functional changes which allow the reader to determine whether to patch now, or defer until the update has been tested. When a CVE has been issued, it also enables the asset owner to look back and identify the window of potential exposure to that vulnerability.

7.3 Improving CPE Quality

As we have discussed in Section 6, the NVD curates the CPE vectors for each CVE, where vendors only can contribute the CVSS

	Mean days (Std. Dev.)	>1 Week	Within 6 Months	Within 1 Year
All Ven- dors	268 (481.54)	1451	950	1131
Siemens	241 (408.63)	282	187	214

Table 6: Average update time (in days) for CVEs, and numbers of CVEs updates after 1 week, within 6 months and within 1 year

score and vector, and the CWE number for a given CVE. This leads to a lack of domain knowledge, for example how products are known within a vendor’s portfolio. Using the Siemens S7-1200 example given earlier in this paper, the S7-1200 range shares common firmware, and is closely related to the SIPLUS S7-1200 range, where SIPLUS products generally receive additional certification for deployments. As such, `simatic_s7-1200` would be an appropriate component in the CPE, as this includes the type of product (SIMATIC is Siemens’ industrial automation portfolio), and the range of products. It also ensures that as new S7-1200 products are introduced, they would be automatically brought into scope given the shared core firmware.

There are two potential solutions to this. The first is that where the vendor is a CNA, they compile a list of CPE vectors based on some agreed internal standard, where the format is well-understood by asset owners. An alternative option is that, when a new product is issued, a base CPE is created by the vendor and supplied to the CPE dictionary with authoritative references. This ensures that if a vulnerability is found to affect that range, the vendor’s accepted representation is used by the NVD assessors.

This does not resolve the issue when additional products are found to be vulnerable, or the issue has been resolved. Currently, the CPE is set and left, where, as an example with a significant update by Siemens in February 2020 to include the SIPLUS ranges in scope of a number of vulnerabilities, the NVD CPE list was never updated, even though the CVE description reflected the change. Moreover, when a vulnerability was patched, the ‘end version’ was never set on the CVE. For an asset owner querying if their asset is vulnerable, they will continue to be informed that all versions are affected. By updating CPE vectors, the CVE is a living record, and an asset owner can be confident they are not vulnerable, as they run a newer firmware than what has been stated as the ‘end version’. Currently, CPEs within a CVE will state the start and end versions, which can be populated once a vulnerability is resolved.

From our observations when reviewing Siemens CVEs, 2 CVEs were updated in July 2020 (CVE-2017-2680 and CVE-2017-2681), where the description was amended, removing the list of affected products. This makes the quality of the CPE listings even more critical, as there is now only one location in which the affected products may be referenced. Whilst this enables a ‘single-source of truth’ approach, it removes the human-readable element for an asset owner.

One question that these solutions raise is around their implementation. Whilst efforts should be undertaken to update previous CVEs, this is a significant effort. Instead, by committing to a window of time, e.g. 2-3 years where major updates have taken place (e.g. new products identified to be updated), this allows the process to be embedded, validated and used, with trust in the CPE vector established. For new CVEs, this should not be a significant undertaking,

as vendors will carry out analyses of affected products, and are in a strong position to recommend appropriate CPEs for customers, aligned to the way their portfolios are represented and customers understand the products.

7.4 Rapid Updates of CVEs as Vendor Updates are Published

Table 6 shows the average number of days taken to update a CVE, and the number of CVEs updated within 6 months and 1 year, based on the last update date. This shows that more than 300 CVEs (for all ICS CVEs) and over 70 for Siemens were updated over a year after publication. As new vulnerable devices are identified based on the updates to vendor advisories observed, we would expect this number to be higher. These update dates, however, do not provide much information about the nature of the update, where NVD provides a change log when a CVE has been updated, either in the MITRE authoritative record, or the NVD assessment changes.

In our review of Siemens' CVEs, we observed updates to the Siemens Security Advisories, some which were updated on the same day as the CVE being published (e.g. CVE-2019-12259). Where Siemens had issued an update, which is cited on the NVD page as a reference, the corresponding description and CPE list had not been updated. This understated the number of vulnerable devices, where the CVSS score assigned was 9.8, making it a critical vulnerability.

It is, therefore, imperative that CVEs are updated as new information comes to light from vendors, as sparsely-distributed content impacts the accessibility of vulnerability information, where the CVE should aggregate all available information and be as accurate and current as possible. This is partly resolved with our proposal of enabling vendors to recommend applicable CPEs, where, as new products are identified, or they confirm products are not vulnerable as originally thought, enabling accurate, current information to be contributed to the CVE. In 2018, Dragos emphasised reviewing related products before a CVE was published to ensure a thorough assessment was conducted and the attack surface well-understood [4], which could have prevented a number of the inaccuracies we identified in our dataset.

8 CONCLUSION

Providing timely and accurate insights to inform asset owner risk management strategies is vital. In this paper, we discover that, in stark contrast to traditional IT vulnerabilities, the timespan of a vulnerability being "in the wild" greatly varies - on average 5.3 years - delays which can have serious safety implications. Additionally, what the CVE says is affected is not always the case, with significant variations, with 15.5% of Siemens CVEs understating affected devices in the list of CPEs. Both these issues critically impact the accuracy and validity of information to asset owners aiming to understand and mitigate risks to their infrastructures, with potentially serious consequences. By improving ICS vulnerability information, asset owners and the supply chain will be better informed when defining their risk management strategies. This, in turn, calls for better support for third-party vulnerability research, further certifications on testing, more control for supply chains to contribute to CVE information reducing this exposure window and, in general, improving the quality of information captured by CVEs.

We are in the process of reporting our findings regarding the inconsistencies within vulnerabilities to both NIST and Siemens.

ACKNOWLEDGMENTS

Funding for this paper was provided by the National Cyber Security Centre UK (NCSC UK), Research Institute in Trustworthy Inter-Connected Cyber-Physical Systems, UK Rail Research and Innovation Network (UKRRIN) and the Engineering and Physical Sciences Research Council (DYPOSIT (EP/N021657/2) and PETRAS2 (EP/S035362/1)).

REFERENCES

- [1] O. Andreeva, S. Gordeychik, G. Gritsai, O. Kochetova, E. Potseluevskaya, S. I. Sidorov, and A. A. Timorin. Industrial Control Systems Vulnerabilities Statistics. *Kaspersky Lab, Report*, 2016.
- [2] Apple Inc. Apple Security Research Device Program. <https://developer.apple.com/programs/security-research-device/>.
- [3] T. Dawson. Who Were the Leading Vendors of Industrial Controls in 2017?, 2017. <https://www.interactanalysis.com/who-were-the-leading-vendors-of-industrial-controls-plcs-and-dcs-in-2017/>.
- [4] Dragos. 2018 Year in Review - Industrial Controls System Vulnerabilities, 2018.
- [5] Dragos. 2019 Year in Review - ICS Vulnerabilities, 2019.
- [6] J. Gardiner, B. Craggs, B. Green, and A. Rashid. Oops I did it again: Further adventures in the land of ICS security testbeds. In *ACM Workshop on Cyber-Physical Systems Security and Privacy (CPS-SPC)*. ACM Press, 2019.
- [7] General Electric. Achilles Communications Certified Products. <https://www.ge.com/digital/applications/achilles-communications-certified-products>.
- [8] General Electric. Achilles System Certification (ASC) from GE Digital: FAQ. https://www.ge.com/digital/sites/default/files/download_assets/Achilles-System-Certification-FAQ.pdf.
- [9] K. E. Hemsley, E. Fisher, et al. History of Industrial Control System cyber incidents. Technical report, Idaho National Lab.(INL), Idaho Falls, ID (United States), 2018.
- [10] Industrial Control Systems Cyber Emergency Response Team. ICS-CERT Annual Assessment Report FY 2016, 2016.
- [11] B. Johnson, D. Caban, M. Krotofil, D. Scali, N. Brubaker, and C. Glycer. Attackers Deploy New ICS Attack Framework "TRITON" and Cause Operational Disruption to Critical Infrastructure, 2017.
- [12] Kaspersky ICS CERT. Threat Landscape for Industrial Automation Systems, 2020.
- [13] National Cyber Security Centre. Cyber Assessment Framework, 2020. https://www.ncsc.gov.uk/files/NCSC_CAF_v3.0%20.pdf.
- [14] V. H. Nguyen and F. Massacci. The (Un)Reliability of NVD Vulnerable Versions Data: an Empirical Experiment on Google Chrome Vulnerabilities. In *Proceedings of the 8th ACM SIGSAC symposium on Information, computer and communications security*, pages 493–498, 2013.
- [15] A. Rashid, J. Gardiner, B. Green, and B. Craggs. Everything is Awesome! Or is it? Cyber Security Risks in Critical Infrastructure. In *Critical Information Infrastructures Security, CRITIS 2019, Linköping, Sweden*. Springer, 2019.
- [16] J. Ruohonen. A look at the time delays in CVSS vulnerability scoring. *Applied Computing and Informatics*, 15(2):129 – 135, 2019.
- [17] W. Schwab and M. Poujol. The State of Industrial Cybersecurity 2018, 2018. <https://ics.kaspersky.com/media/2018-Kaspersky-ICS-Whitepaper.pdf>.
- [18] M. Shukla, S. D. Johnson, and P. Jones. Does the NIS implementation strategy effectively address cyber security risks in the UK? In *2019 International Conference on Cyber Security and Protection of Digital Services (Cyber Security)*, pages 1–11. IEEE, 2019.
- [19] Siemens. Siemens Vulnerability Handling and Disclosure Process. <https://new.siemens.com/global/en/products/services/cert/vulnerability-process.html>. Accessed: 2020-07-12.
- [20] R. J. Thomas and T. Chothia. Learning from Vulnerabilities - Categorising, Understanding and Detecting Weaknesses in Industrial Control Systems. In *Computer Security*, Cham, 2020. Springer International Publishing.
- [21] J. Walden, J. Stuckman, and R. Scandariato. Predicting vulnerable components: Software metrics vs text mining. In *2014 IEEE 25th International Symposium on Software Reliability Engineering*, pages 23–33, 2014.
- [22] T. Wallis and C. Johnson. Implementing the NIS Directive, driving cybersecurity improvements for Essential Services. In *2020 International Conference on Cyber Situational Awareness, Data Analytics and Assessment (CyberSA)*, 2020.
- [23] B. Wang, X. Li, L. P. de Aguiar, D. S. Menasche, and Z. Shafiq. Characterizing and Modeling Patching Practices of Industrial Control Systems. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 1(1):1–23, 2017.