

A Quantitative Study of Vulnerabilities in the Internet of Medical Things

Hervé Debar¹^a, Razvan Beuran² and Yasuo Tan²

¹*SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France*

²*School of Information Science, Japan Advanced Institute of Science and Technology, Nomi, Japan*

Keywords: Cybersecurity, Medical Internet of Things.

Abstract: Medical objects, small or large, increasingly rely on digital technologies to monitor patients or deliver care. They form a part of our digital critical infrastructure, that can be significantly impacted by cyberattacks. For example, the Wannacry ransomware shut down hospitals in Europe for hours, even days. This paper analyzes recent vulnerabilities that have affected medical objects, and present findings related to the characteristics of these vulnerabilities. It will then use these findings to propose ideas for improved cybersecurity in the medical IoT. One of the key findings of the paper is that it demonstrates the effect of regulations enacted worldwide in early 2017, requiring critical infrastructure operators and providers to proactively publish information about vulnerabilities.

1 INTRODUCTION

Medical objects, small or large, increasingly rely on digital technologies to monitor patients or deliver care. They form a part of our digital critical infrastructure, that can be significantly impacted by cyberattacks. For example, the Wannacry ransomware shut down hospitals in the UK for hours, even days (Martin et al., 2017). Granted, other types of critical infrastructures were impacted elsewhere in the world by this same threat. However, this incident shows that a significant number of healthcare processes, forming the Internet of Medical Things, are relying on information and communication technologies.

Medical infrastructures form part of our Critical Infrastructure, or essential services. As such, recent regulations have introduced an obligation for operators of critical infrastructures to disclose incidents, and an obligation for vendors of products used in these infrastructures to publish vulnerability information. An example of sites providing this information is the US-based ICS CERT advisories (US Department of Homeland Security, 2019).

Using ICS CERT advisories, this paper focuses on advisories that apply to the “Healthcare and Public Health” critical infrastructure sector, identified by the “ICSMA” identifier in the advisory organization. Compared with other sources of information (e.g. articles in journals), this information is reliably struc-

ured, up to a point where it is possible to envision further work to analyze this information.

This paper analyzes recent vulnerabilities that have affected medical objects, and present findings related to the characteristics of these vulnerabilities. It will then use these findings to propose ideas for improved cybersecurity in the medical IoT, leveraging the conceptual description of the NIST cybersecurity framework.

2 CONTEXT AND RELATED WORK

This work should be understood in the context of critical infrastructure protection, at a time where regulations such as the Network and Information Systems security (NIS) directive (NIS Directive, 2016) impact the way industries labelled as critical infrastructures, and the vendors which serve these industries, must disclose vulnerability information to the regulators.

Medical environments are critical infrastructures, and providing good cybersecurity measures may prove essential to patient health (Martin et al., 2017) already now. Medical environments have specific requirements, however, which have led to specific solutions being developed in the past such as *break the glass* access control (Marinovic et al., 2011). One of our objectives is to evaluate if medical devices have a different vulnerability profile than what we are classi-

^a <https://orcid.org/0000-0002-1344-4167>

cally used to in information and communication technologies, and if we can quantify this difference.

There have been a few studies of cybersecurity risk in medical environments. Fu and Blum (Fu and Blum, 2013) analyze cybersecurity reports of medical device manufacturers to the Food and Drug Administration (FDA). They provide a qualitative analysis indicating that there are significant cybersecurity vulnerabilities related to integrity and availability, but do not quantify this information any further.

Kramer and Fu (Kramer and Fu, 2017) study the specific case of connected pacemakers by analyzing the advisory provided by the Food and Drug Administration (US). It concludes that the information provided in the advisory could be improved. Our root source is different, as we look at the data from a more global perspective, and information we process spans multiple vendors. Also, the FDA is concerned primarily with patient safety and is producing recommendations to this effect, whereas we are analyzing recommendations that apply to the command and control ICT infrastructure in the medical sector.

Kruse *et al.* (Kruse *et al.*, 2017) analyze publications from the medical databases. They extract 31 documents related to cybersecurity in the medical sector. It highlights that cybersecurity is a significant issue for the medical sector, and that this is reflected in regulations such as HIPAA (Annas *et al.*, 2003). However, it does not provide tools or quantification of the issues they highlight.

Coventry and Branley (Coventry and Branley, 2018) also analyze the medical scientific literature to understand why healthcare environments are vulnerable to cybersecurity issues. They elaborate that there are an increasing number of connected devices per patient (15 is mentioned in their study) and that healthcare environments host valuable data. However, they review specific incidents but do not provide a global view of vulnerabilities.

As a summary, the existing literature focuses on qualitative analysis, while we attempt to quantify the issues. Furthermore, many of these studies are driven by the medical world, while we attempt to bring an ICT perspective to our analysis.

The analysis heavily relies on structured information provided by the Common Weakness Enumeration (CWE) (Martin, 2007) and the Common Vulnerability Scoring System (CVSS) (Mell *et al.*, 2006). As a by-product of this study, we also expect to show how effective or useful these widely-used schemes are to understand vulnerabilities.

3 METHODOLOGY

The objective of the paper is to systematically analyze advisories of the medical critical infrastructure, as documented by the ICS CERT advisories. The following methodology was created to support this systematic analysis.

Collecting and Parsing Advisories. Advisories have been collected from the ICS-CERT website and separated into structured information. This is performed based on the structure of the advisory as shown on the website.

Enhancing Advisories. Advisories are enhanced in several specific ways, to get information from the Common Weakness Enumeration, patch status and other vendor information.

Queries for Analysis. Once the dataset is structured and clean, we ran a number of queries to get answers to questions such as complexity of advisories, availability of patches, time between discovery and patch, complexity of the attack vector.

3.1 Collecting and Parsing Advisories

An advisory is a well formatted document, consisting of four or five sections. For each advisory, we use the identifier and title provided by CERT, as well as the date of first publication.

The first section, “executive summary”, is a bullet list of summary information providing information about the CVSS score, the attack vector and skill level required, the vendors involved, the products involved and the type of vulnerability. As this is a summary information, the technical details described hereafter are more precise, particularly on products and versions affected.

The second (optional) section “update information” contains update information. As advisories are revised when additional information becomes available, this section contains time-related information about the update. The text of the advisory, and in particular the technical details and the mitigation section contain revision marks that indicate the update. This section is not considered in the analysis, as we only address the most up to date (at the time of collection, January 2019) version of the advisory.

The third section, “Risk analysis”, is a short sentence about risk related to the advisory as a whole. As risk is more precisely addressed by the CVSS vector, this section is not considered in the analysis.

The fourth section, “technical details”, is the most detailed content of the advisory. It is composed of two subsections. Subsection “affected products” lists with

significant details the product names and versions that are affected by the advisory. This section also provides information about the products that are not affected by the advisory. The second subsection, “vulnerability overview”, provides a detailed list of each of the vulnerabilities that are affecting the products. Each vulnerability (CVE entry) is reported in an individual subsection, whose title references the CWE entry categorizing the CVE information.

Each vulnerability description is thus included in a third-level subsection. The title of the vulnerability subsection references a CWE category of vulnerability (text, identifier and link). The content of the description includes a few lines of text describing in more details the mechanism of the vulnerability and its consequences. The last sentence of each section provides a CVE identifier, a CVSS v3 score and a CVSS v3 vector.

The final section, “Mitigations”, provides information to remediate the vulnerability. It is informally split in two parts. The first provides information from the vendor about availability (or not) of patches, and what they recommend to address with respect to the vulnerability. This part also provide links to additional vendor information on their web site. The second part is a standard text provided by ICS-CERT on how to generally mitigate vulnerabilities in critical infrastructures.

3.2 Enhancing Advisories

Once the initial information is collected, further processing is required to enhance and normalize the information collected. This enhancement process relies on third party information that is directly linked to by the advisories.

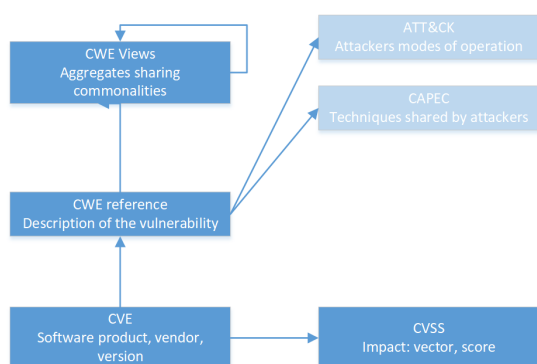


Figure 1: Structuration of vulnerability information.

As shown in figure 1, the most basic vulnerability directory is CVE, directly associated with a CVSS vector. However, the *technical details* section of advi-

sories is sub-structured using the CWE structure, and the CVE/CVSS information is given only at the end of each subsection. This means that the number of CWE in an advisory is lower or equal to the number of CVE it reports.

The first enhancement relates to the completion of the CWE information associated with each advisory. CWE contains roughly 1000 entries, classified as base information, or as class. A class regroups vulnerability types that share certain characteristics. For our purposes, even regrouping CWE provides too segmented information. Therefore, we are using also the “Software Fault Patterns” (SFP) view of CWE, which provides a typology of frequent errors made by developers. This view has two layers, a very compact aggregate of 10 high-level common software mistakes (referred to as the SFP1 view), and a second layer refining some of these high-level clusters in more precise sub-clusters.

As shown in figure 1, there are other substructures available, such as CAPEC and ATT&CK. This analysis is left for future work.

The second enhancement is the segmentation of the CVSS vector. We transformed all CVSSv2 information included in the older advisories in the CVSSv3 format using a fixed translation scheme. We then segment the CVSS vector into its eight components, to ensure that we are able to analyze in details the attack vectors and the impact.

Concerning date and time, we take specific care to extract the year of both the CVE and the advisory titles. This provides us with a timeline representing roughly the beginning and the end of the vulnerability spectrum. CVE entries are reserved when a researcher suspects a vulnerability and requests an entry in the CVE repository from MITRE. Thus, the date in the CVE entry represents roughly the time of discovery. Advisories are generally published when vendors have had time to confirm the existence of the vulnerability and provide a solution.

3.3 Completing Advisories with Device Type

Certain fields must be manually analyzed in order to ensure accurate and complete information. We perform manual completion of advisories on two counts, understanding what kind of product is vulnerable, and what kind of mitigation solution is available.

The first manual analysis performed is linked to the understanding of what the product is, and more precisely determining if the advisory describes a device vulnerability (thus a vulnerability in the embedded software or firmware), or if it is describing an

application vulnerability. ICS-CERT advisories address critical infrastructure, but not necessarily devices. Medical software platforms such as imaging management platforms are purely software, but they form part of the critical infrastructure and as such are subject to advisories. So manually reading the vendor documentation about the product is required to determine whether the vulnerable product is a medical device, or pure software.

Traditional industrial control systems differentiate between *sensors* (temperature, pressure, etc.), which provide input data to the control processes, and *actuators* (valves, motors, etc.) which apply changes to the control system. In medical devices, we have a need for a finer grain classification, as sensors such as imaging devices need to subject the body to physical effects, typically radiation, to obtain the desired measurement. We thus will use the following four classes for devices, extending the transducer capabilities of the taxonomy proposed in NISTIR 8228 (Boeckl et al., 2019).

PassiveSensing. *The ability to Provide an Observation of an Aspect of the Physical World in the Form of Measurement Data* (Boeckl et al., 2019). In the context of this study, we restrict devices in this category to measurements that cannot have a negative effect on humans. Examples include temperature or blood pressure.

Actuating. *The Ability to Change Something in the Physical World* (Boeckl et al., 2019). In the healthcare domain, the environment might be the patient, but also the operator of the device. Devices in this category directly inject treatment to the patient. This category includes for example connected syringes, pacemakers or insulin pumps.

ActiveSensing. This extension of NISTIR 8228 covers the case of a device which requires some actuation to do the sensing. In the healthcare domain, this might be for example the case of an imaging device that needs to emit X-rays to capture the image. There might be an adverse effect of X-Rays to the patient or the operator. The reason for having a separate category is that there are recorded incidents where patients were subjected to inappropriate levels of radiations by operators, generally by accident.

IndirectActuating. This extension of NISTIR 8228 covers the case of a device management platform, such as an aggregator or a programming tool. It is not a direct device, but it is so close to the device that an attack on it might influence the device itself. An example in the healthcare domain includes syringe management platforms, drug de-

livery platforms.

3.4 Completing Advisories with Patch Information

The second aspect of advisories that requires manual analysis is the understanding of the possible mitigation solutions associated with the advisory. The “mitigation” section of advisories is much less structured than the other parts. Furthermore, it frequently links to the security section of the vendor website, where further interactions are needed to find which piece of vendor-provided information actually applies. Thus, there are 3 situations that we are confronted with:

1. **Unknown:** We have not been able to extract information relative to the existence of a patch.
2. **No Solution Provided:** The text clearly indicates that the vendor will not fix the vulnerability. The vendor does not provide a robust solution that enable the user to continue using the device with trust.
3. **Solution Provided:** The text clearly indicates that the vendor has fixed or will fix the vulnerability.

To exhaustively describe the different situations encountered in advisories, we define the following six patch status:

Unknown. No information could be extracted relative to the existence of a solution provided by the vendor.

No The status ‘No’ indicates that no solution will be provided by the vendor.

Upgrade. The status ‘Upgrade’ indicates that the vulnerability is patched in newer versions of the product; the vulnerable version is out of date and the user should buy a newer version.

Mitigation. The status ‘Mitigation’ indicates that a workaround is described by the vendor. The vulnerability remains in the product, but exploitation is harder or effects are mitigated.

Announced. The status ‘Announced’ indicates that the vendor has given a date and methodology for the release of the patch.

Yes The status ‘Yes’ indicates that the solution has been published and is available, and can be deployed by the users.

The favorable situation is clearly when a patch is either announced or available. In the other cases, the end-user may have to cover significant expenses renewing his product or changing its modes of operation to solve the issue.

3.5 Remaining Issues

The following additional enhancements could be applied to the dataset, in order to provide increased quality.

Certain CWE, particularly the ones associated with older vulnerabilities, are not precise enough. This is illustrated for example by the fact that certain CWE are tagged as “used by NVD before 2016”. They could be replaced by more precise ones (not done yet), focusing particularly on categories that provide an aggregated view but not a precise idea of the problem. The text associated with the description of the vulnerability could be mined to specify more precisely the issue. One example of this is CWE-310, which associated with *legacy SSL* would be more accurately described by CWE-327, whereas CWE-310 associated with *certificate* would be more accurately described by CWE-295.

There is a need to strengthen the process with respect to products and vendors. For example, there is an issue related to mergers and acquisitions between Becton and Carefusion, the later bought by the former. Hence there is a product with the same name, but different versions, which is attached to two vendors.

There is a need to specialize the confidentiality impact and to adapt it to the medical domain. For example, the process needs to determine if Personally Identifiable Information (PII) and Personal Health Information (PHI) is impacted. We assume that financial information will never be impacted in the medical systems, as it resides purely in ICT service platforms.

3.6 Analyzing Advisories

At this stage, we obtain a clean dataset of advisories, associated with CWE records. In the end, the dataset consists of 55 fully complete advisories, referencing as a whole 204 CVE entries. The dataset references 62 products, 48 devices in one of the four categories listed in section 3.3, and 14 referencing vulnerabilities in software platforms. 40 of the 55 advisories are related to devices.

Out of the 24 vendors, 20 are device vendors and 6 are software vendors. Therefore, two vendors (Becton Dickinson and Philips) are present in the data set with vulnerabilities both on software and on devices. Out of the 24 vendors, only 6 had reports related to several years. Philips reported during 4 years, Becton Dickinson 3. This tends to indicate an increase in awareness and information sharing. 11 device vendors reported an issue in 2018.

As shown in Figure 2, the importance of reporting

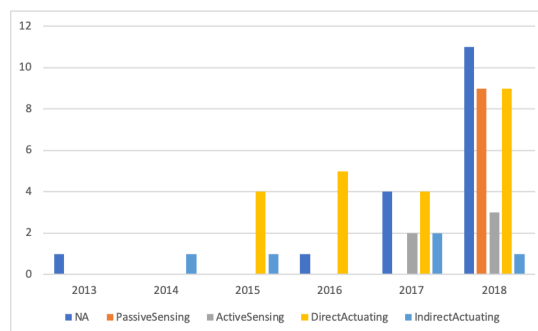


Figure 2: Number of advisories per year per device type.

is improving significantly over the years. All categories of devices are providing information in 2018. The category of PassiveSensing is also emerging in 2018. Since the requirements to report cybersecurity issues in the critical infrastructure sector have been enacted, there has been a significant number of advisories reported.

4 ANALYSIS RESULTS

4.1 Time Lag between CVE and Advisory

As is well known in the software world, there might occur a significant time lapse between the discovery of a vulnerability and the actual publication of the advisory. So we would like to get a quantitative answer for the question of *How long does it take for a vulnerability to be fixed ?*

This time lapse in the dataset can be modeled by the difference between the CVE year and the year of publication of the advisory. The former represents an approximation of the time at which the discovery of the vulnerability is made by a white hat hacker, while the publication of the advisory indicates that the vendor has been alerted and has handled the problem, often in recent years by providing a patch.

Table 1 counts the number of vulnerabilities having as year of the advisory the line value (from 2013 to 2018) and as year of the CVE the column value (from 1999 to 2018). The expectation is that if vulnerabilities are discovered and handled by vendors rapidly, we should obtain a diagonal table.

The most reassuring result is that, at the year granularity (which is quite large), many vulnerabilities are discovered and disclosed either the same year, or the year after, and we obtain a diagonal, as expected. However, there are a few outliers that required further analysis.

Table 1: Difference in years between the year of the CVE and the year of the advisory.

Year Adv.	Year as indicated in the CVE label (CVE-YEAR-NNNN)																			
	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018
2013															1					
2014																4				
2015								1								9	17			
2016																	1	8		
2017										1							4	2	16	
2018	1		1	2	1	2	1			1		1	5	4	4	4		1	59	65

The CVE-1999 entry in fact references the *ping-of-death vulnerability* and could be considered an outlier. The actual date of the ping-of-death vulnerability for windows platforms is 1996, and for BSD-Unix systems around 1985. Therefore, while it is surprising that a product analyzed in 2018 should exhibit such an old vulnerability, the time difference between CVE and advisory is justified.

Another outlier is the ICSMA-18-037-02 vulnerability, also addressed in section 4.2. While the advisory itself is from 2018, the vulnerability affects many products of the same family, which have been developed and retired over the years. Hence, the wide spread in CVE within the same advisory. This explains most of the time lag (27 out of 30 entries).

As a conclusion, while there is clearly a greater time lapse than the best practice of 90 days between private disclosure to vendor and public disclosure, there is a significant effort to tackle vulnerabilities in the critical healthcare infrastructure within a reasonable amount of time.

4.2 Complexity of Advisories

Another question is the complexity of advisories. If an advisory contains many vulnerabilities, it is likely to be both more significant and more difficult to process for critical infrastructure operators.

Advisories for software report in general only one CVE per advisory. Our data contains only one outlier reporting 35 vulnerabilities (ICSMA-18-058-02), related to a web portal. This might be considered as representative of a large software package patching many vulnerabilities related to the underlying operating system.

Figure 3 counts the number of CVE entries in advisories for medical devices. It shows that the advisories tend to cover multiple CVE references. 11 advisories cover a single CVE entry (out of 55), while 14 contain 2 CVE references, and another 14 contain between 3 and 8 CVE references. So while advisories are more complex for devices than for products, the complexity remains limited. There are two outliers, reporting 15 and 27 CVE entries for a single advisory.

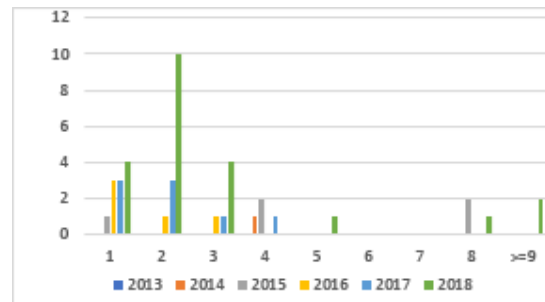


Figure 3: Number of CVE entries per advisory for devices.

They cover two different cases.

ICSMA-18-107-02 reports many vulnerabilities related to the use in the device of the MS operating system. In the dataset, this is artificial in the sense that the advisory itself does not list any of the vulnerabilities, but these can be found in the vendor information associated with the advisory. This reflects the fact that at least this vendor ignored the lifecycle of the operating system it used in its device.

ICSMA-18-037-02 reports 27 CVE entries for the same CWE (CWE 287 – Improper authentication, likely the use of hard coded password) in a family of products over the years. The CVE entries range from 2001 to 2017, indicating that the problem has been known for a significant duration, but was not addressed by the vendor until very late. The segmentation of CVE is also very small and might be hard to address.

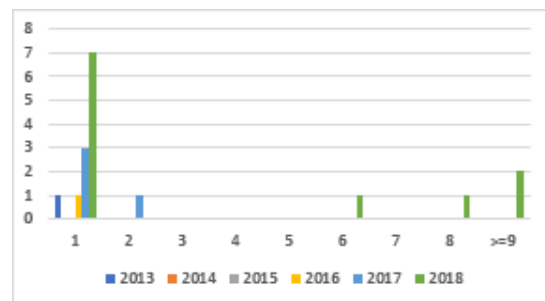


Figure 4: Number of CWE entries per advisory for devices.

Figure 4 depicts the same information but for

CWE entries. It shows that for example the two outliers presented above have disappeared, because both of them are related to only a few types of vulnerability. In particular, ICSMA-18-037-02 has completely disappeared in the figure.

This brings two comments. First, it confirms that CWE is an effective description of the status of vulnerabilities. Second, it indicates that medical devices may suffer from a single (or two) type of flaws, even if there are multiple CVE referenced. This may indicate that there is hope for more efficient mitigation.

Another aspect of advisory complexity is related to CVSS scores. The profile of severities remains relatively stable over the years. It is impossible to consider that there is an increase in risk related to an increase in severity. However, there is no significant decrease either, meaning that vulnerabilities continue to be introduced at a significant rate.

4.3 Analysis of Common Weakness Enumeration

Common Weakness Enumeration is a taxonomy of vulnerabilities that is frequently included in the advisories. The following analysis attempts to understand the type of fault that is really occurring, based on CWE references that are associated with many entries. This part of the analysis uses as pivot the CWE references. There are 63 different CWE references in the database. They are present in a large portion of the dataset. Only one advisory for devices does not properly reference a CWE entry.

Out of these 63 references, 2 are compounds concerning composites, CWE-352 and CWE-384. They mean that several basic flaws must be present for the vulnerability to exist. These are difficult to classify further, although CSRF is a well-known OWASP issue.

Six other CWE references are categories. These entries group several CWE entries that share a common characteristic. It probably means that the identification of the vulnerability was not precise enough. 1 outlier is specified as graph, which has the same characteristics. In terms of structure, they are classified as incomplete and are excluded from further analysis.

Table 2 presents the list of simple CWE entries for devices in the dataset, with their name, the frequency of apparition overall and in 2018, and the first-level software fault pattern cluster to which they belong.

The first result from the query is that the most present issue, by far, is related to authentication. CWE-287 is essentially a failure to validate fully the authentication data submitted to the device. A similar issue is CWE-345. while CWE-259 and CWE-

798 relate to the fact that authentication information is hard-coded and cannot be modified. The same issue appears with CWE-321, but the use of cryptography indicates an increased awareness of security and hence is less frequent than password.

The most repeated issue is related to *improper authentication* (i.e. the absence of a mechanism to control access to the device), closely followed by *default password*. These issues, which have completely disappeared in traditional software development, may indicate that devices are developed by non-security specialists or do not follow risk assessment procedures.

Several of the Information Leak category entries also indicate leakage of authentication information.

This situation remains in 2018. The main difference between the overall data and the 2018 focus seems to be the removal of cryptographic errors, which may indicate a better maturity of the developers in using encryption.

CWE-323 appears only once in the dataset. This references the infamous Key Reinstallation Attack (KRACK) vulnerability in WPA2. Since networking capabilities are clearly part of modern devices, and one of their attack channels, it is surprising that there are not more vendors reporting this. This probably means that there are other advisories to come, or that this vulnerability is silently patched by vendors.

The same information for software brings up the traditional issues of tainted input and memory management, which also exist in devices but less frequently.

4.4 Analysis of the CVSS Vector

This section focuses on analyzing two subparts of the CVSS vector, the *Attack vector* and the *Attack complexity*. The attack vector indicates the capability of the attacker to carry out its attack remotely (indicated by adjacent or Network in table 3). The attack complexity indicates the difficulty that the attacker has in carrying out the attack successfully and is qualified by either low or high. The table order indicates decreasing number of vulnerabilities.

As unfortunately expected, vulnerabilities are widely exploitable over the open Internet with a low attack complexity, as indicated by the last line of table 3. In that respect, devices tend to exhibit more failings than pure software (where physical access is considered irrelevant). More worrisome, out of the 42 network-accessible, low-complexity vulnerabilities that were reported in 2018, 41 also require no particular privileges to be exploited. This may indicate that there has been no significant progress in protecting devices that are in many cases directly connected

Table 2: Count of CWE entries in the dataset and in 2018.

CWE-Ref	CWE title	Count	2018	SFP Primary
CWE-120	BUFFER COPY WITHOUT CHECKING SIZE OF INPUT	1	1	Memory Access
CWE-125	OUT-OF-BOUNDS READ	1	1	Memory Access
CWE-200	INFORMATION EXPOSURE	1	1	Information Leak
CWE-23	RELATIVE PATH TRAVERSAL	1	1	Path Resolution
CWE-250	EXECUTION WITH UNNECESSARY PRIVILEGES	1	1	Privilege
CWE-254	SECURITY FEATURES	1	0	NA
CWE-256	UNPROTECTED STORAGE OF CREDENTIALS	1	1	Information Leak
CWE-264	PERMISSIONS PRIVILEGES AND ACCESS CONTROLS	1	0	NA
CWE-295	IMPROPER CERTIFICATE VALIDATION	1	1	NA
CWE-323	REUSING A NoNCE KEY PAIR IN ENCRYPTION	1	1	Cryptography
CWE-330	USE OF INSUFFICIENTLY RANDOM VALUES	1	0	Predictability
CWE-377	INSECURE TEMPORARY FILE	1	0	Information Leak
CWE-427	UNCONTROLLED SEARCH PATH ELEMENT	1	1	Tainted Input
CWE-434	UNRESTRICTED UPLOAD OF FILE WITH DANGEROUS TYPE	1	1	NA
CWE-460	IMPROPER CLEANUP ON THROWN EXCEPTION	1	1	Exception Management
CWE-601	URL REDIRECTION TO UNTRUSTED SITE AKA OPEN REDIRECT	1	0	Tainted Input
CWE-668	EXPOSURE OF RESOURCE TO WRONG SPHERE	1	1	Information Leak
CWE-693	PROTECTION MECHANISM FAILURE	1	1	Other
CWE-732	INCORRECT PERMISSION ASSIGNMENT FOR CRITICAL RESOURCE	1	0	Access Control
CWE-755	IMPROPER HANDLING OF EXCEPTIONAL CONDITIONS	1	1	Exception Management
CWE-923	IMPROPER RESTRICTION OF COMMUNICATION CHANNEL TO INTENDED ENDPOINTS	1	1	NA
CWE-119	IMPROPER RESTRICTION OF OPERATIONS WITHIN THE BOUNDS OF A MEMORY BUFFER	2	0	Memory Access
CWE-20	IMPROPER INPUT VALIDATION	2	2	Tainted Input
CWE-257	STORING PASSWORDS IN A RECOVERABLE FORMAT	2	2	Information Leak
CWE-260	PASSWORD IN CONFIGURATION FILE	2	1	Information Leak
CWE-285	IMPROPER AUTHORIZATION	2	0	Access Control
CWE-294	AUTHENTICATION BYPASS BY CAPTURE-REPLAY	2	1	Channel
CWE-300	CHANNEL ACCESSIBLE BY NoNENDPOINT AKA MAN-IN-THE-MIDDLE	2	0	Channel
CWE-319	CLEARTEXT TRANSMISSION OF SENSITIVE INFORMATION	2	1	Information Leak
CWE-320	KEY MANAGEMENT ERRORS	2	0	NA
CWE-321	USE OF HARD-CODED CRYPTOGRAPHIC KEY	2	0	Authentication
CWE-356	PRODUCT UI DOES NoT WARN USER OF UNSAFE ACTIONS	2	2	UI
CWE-400	UNCONTROLLED RESOURCE CONSUMPTION AKA RESOURCE EXHAUSTION	2	0	Resource Management
CWE-749	EXPOSED DANGEROUS METHOD OR FUNCTION	2	1	Other
CWE-78	IMPROPER NEUTRALIZATION AKA OS COMMAND INJECTION	2	2	Tainted Input
CWE-920	IMPROPER RESTRICTION OF POWER CONSUMPTION	2	1	NA
CWE-312	CLEARTEXT STORAGE OF SENSITIVE INFORMATION	3	0	Information Leak
CWE-345	INSUFFICIENT VERIFICATION OF DATA AUTHENTICITY	3	1	Authentication
CWE-522	INSUFFICIENTLY PROTECTED CREDENTIALS	3	1	Information Leak
CWE-121	STACK-BASED BUFFER OVERFLOW	4	1	Memory Access
CWE-284	IMPROPER ACCESS CONTROL	4	4	Access Control
CWE-311	MISSING ENCRYPTION OF SENSITIVE DATA	4	2	Information Leak
CWE-94	IMPROPER CONTROL OF GENERATION OF CODE AKA CODE INJECTION	4	0	Tainted Input
CWE-798	USE OF HARD-CODED CREDENTIALS	6	4	NA
CWE-259	USE OF HARD-CODED PASSWORD	8	2	Authentication
CWE-287	IMPROPER AUTHENTICATION	9	6	Authentication

Table 3: CVSS vector comparison between devices and software, in the dataset and in 2018.

Attack characteristics		Complete dataset		Year 2018	
Vector	Complexity	Devices	Software	Devices	Software
Physical	High	7	0	5	0
	Low	5	0	2	0
Local	High	6	0	6	0
	Low	13	6	8	5
Adjacent	High	20	0	17	0
	Low	5	3	4	3
Network	High	12	26	7	26
	Low	67	30	42	25

to the Internet.

A few attacks do require physical access to the device, which in healthcare settings is not difficult. This type of vulnerability might reveal difficulties in se-

Table 4: CVSS impact comparison between devices and software in the dataset.

		Confidentiality	Integrity	Availability
Devices	None	14 (10%)	32 (24%)	35 (26%)
	Low	24 (18%)	22 (16%)	19 (14%)
	High	97 (72%)	81 (60%)	81 (60%)
Software	None	12 (18%)	28 (43%)	27 (42%)
	Low	12 (18%)	10 (15%)	6 (4%)
	High	41 (63%)	27 (42%)	32 (24%)

curing devices from the persons who have physical access. In this respect, devices that are not owned by their main users may require specific security properties.

Table 4 extracts from the CVSS vector the usual 3 impacts, Confidentiality, Integrity and Availability

(CIA). The major impact of vulnerabilities on devices is on confidentiality, as only 10% list no impact and 72% list a high impact. Integrity and availability are slightly less impacted. Also, one has to note that the impact of vulnerabilities on software is generally less and less widespread than the impact of vulnerabilities on devices. Half of the device vulnerabilities list a high impact in the 3 dimensions, while this is the case for only 32% of the pure software vulnerabilities.

4.5 Patch Availability

Another interesting question is whether vendors provide solutions for security vulnerabilities. Table 5 indicates the patch information associated with vulnerabilities (CVE entries), according to the nomenclature presented in section 3.4.

Table 5: Patch information for each CVE, for the whole dataset and the years 2017-2018.

Patch Status	All years		Years 2017-2018	
	Devices	Software	Device	Software
Unknown	28 (20%)	2 (3%)	2 (3%)	2 (3%)
No	8 (6%)	1 (1%)	6 (8%)	0 (0%)
Upgrade	22 (15%)	0 (0%)	3 (4%)	0 (0%)
Mitigation	21 (15%)	1 (1%)	10 (14%)	1 (2%)
Announced	14 (10%)	44 (64%)	10 (14%)	38 (63%)
Yes	50 (35%)	21 (30%)	43 (58%)	19 (32%)

The following table depicts a picture that is significantly different between pure software and medical devices. While 94% of the software vulnerabilities are patched, only 45% of them are patched in devices. Vendors rely more heavily on mitigation and upgrade (30%). Also interesting is the fact that for many CVE entries, it is difficult to ascertain the status of the patch.

Fortunately, the difference is reduced if the data is focused on the most recent years 2017-2018. Focusing on this most recent part of the dataset, the proportion of devices effectively protected increases to over 70%.

The devices have a different patching profile than pure software. On one hand this is not surprising because devices require more effort to patch. On the other hand, medical devices (contrary to the ones found in for example smart homes) are managed by professional, and contrary to other settings (e.g. industry) it is possible to have shorter usage lifecycle that could fit a patching model. Furthermore, critical infrastructure has a requirement to maintain these devices in order to remediate cybersecurity issues. The gap is closing in recent years, but effort is still required to continue in this direction.

5 RELATION TO THE NIST CYBERSECURITY FRAMEWORK

The NIST Cybersecurity framework (Shen, 2014) provides a reference for improving critical infrastructure security. The framework core defines five areas (identify, protect, detect, respond and recover) and for each of these areas defines a number of categories and sub-categories where action should be taken to appropriately address the area. Standards cybersecurity controls are associated with each of the sub-categories, to facilitate implementation.

In the context of this work, section 5 attempts to provide recommendations based on the statistics extracted from the data, and other considerations from the literature, for each of the functions.

The recover function covers organizational and communication actions and is outside the scope of the work.

5.1 The Identify Function

The activities in the Identify Function are foundational for effective use of the Framework. Understanding the business context, the resources that support critical functions, and the related cybersecurity risks enables an organization to focus and prioritize its efforts, consistent with its risk management strategy and business needs. Examples of outcome Categories within this Function include: Asset Management; Business Environment; Governance; Risk Assessment; and Risk Management Strategy.

The first category of the identify function is Asset Management (ID.AM): The data, personnel, devices, systems, and facilities that enable the organization to achieve business purposes are identified and managed consistent with their relative importance to business objectives and the organization’s risk strategy. In the context of IoMT, the link in the critical infrastructure between the objects and the management platforms (e.g. image management or device programming) is clearly established through the qualification of advisories. The specificities of the healthcare sector must be reflected in advisories and in the CVSS vector, such as the real possibility of physical harm to patients and operators if the devices are misused, and the requirements to operate in emergency situations. This leads to difficulties as communications and data flows are not necessarily stable, maybe not sufficiently to enable static definition. In our study, this is qualified by the type of device, as described in section 3.3.

The second category of the identify function is Business Environment (ID.BE): The organization’s

mission, objectives, stakeholders, and activities are understood and prioritized; this information is used to inform cybersecurity roles, responsibilities, and risk management decisions.

The advisories studied in this document demonstrate that some of the devices and software suppliers have complied with regulations, at least in the US. However, the number of reported issues remains small overall and further studies should confirm the level of compliance of the healthcare industry as a whole.

The next three categories cover aspects related to governance (ID.GV), Risk Assessment (ID.RA) and Risk Management strategies (ID.RM). They are only covered through the analysis of the CVSS vector. The granularity of the CVSS vector is insufficient to ensure proper satisfaction of these 3 categories, particularly with respect to privacy risks. Mitigation is also insufficiently covered in advisories.

5.2 The Protect Function

The Protect Function supports the ability to limit or contain the impact of a potential cybersecurity event. Examples of outcome Categories within this Function include: Identity Management and Access Control; Awareness and Training; Data Security; Information Protection Processes and Procedures; Maintenance; and Protective Technology. In the context of this work, the analysis of vulnerabilities is the most appropriate source of information to effectively evaluate the recommendations provided by the framework.

The first category of the Protect function is Identity Management, Authentication and Access Control (PR.AC): Access to physical and logical assets and associated facilities is limited to authorized users, processes, and devices, and is managed consistent with the assessed risk of unauthorized access to authorized activities and transactions.

This document demonstrates that authentication issues are by far the most critical vulnerability of medical devices. While physical access is important in the protection of devices, the document demonstrates that devices and software are remotely vulnerable and that physical access may not necessarily be the first required ICT protection, not beyond what is usually carried out in hospital environments. Mitigation actions proposed by vendors unfortunately do not rely on permissions management.

The third category of the Protect function is Data Security (PR.DS): Information and records (data) are managed consistent with the organization's risk strategy to protect the confidentiality, integrity, and availability of information.

Section 4.4 shows that the major impact of vulner-

abilities is on data. However, it does not allow at this stage a differentiation between data at rest and in transit. Further work is required to assess exactly what kind of data is impacted. Overall, the complexity of exploiting vulnerabilities is rather low, indicating failure in development processes for devices.

The other categories of the protect function are less relevant to this work.

Our work demonstrates that the advisories studied in the document are very relevant to support the Protect function. Aspects related to network access control are frequently mentioned in mitigation and in protection, even though they might be hard to realize in practice. Aspects related to access control and identity and access management are frequent causes for vulnerabilities in devices, much more than software which suffers from classic input/output sanitization issues. Data protection aspects, both confidentiality and integrity, are particularly important in the healthcare domain. The study shows that confidentiality in particular is a frequent issue in vulnerabilities.

5.3 The Detect Function

The Detect Function enables timely discovery of cybersecurity events. Examples of outcome Categories within this Function include: Anomalies and Events; Security Continuous Monitoring; and Detection Processes.

The first category of the detect function is Anomalies and Events (DE.AE): Anomalous activity is detected in a timely manner and the potential impact of events is understood.

In the IoMT, communication and data flows may not be stable, as medical environments need to react in emergency situations. The deployment of home care might also require more study on this, especially since many vulnerabilities are remotely actionable. Since many vulnerabilities are related to authentication issues such as default passwords and credentials (section 3.3), attacks may be very similar to regular activity and it might be extremely difficult to deploy misuse or anomaly detection methods. The vulnerabilities studied in this document, either for device or software, do not require complex event processing or multi event processing for the most part. The victim is generally clearly identified and is a single component.

The study provides some worthwhile elements for the Detect function. The main finding is that the vulnerabilities reported for devices touch authentication issues such as default credentials. Attacks exploiting these vulnerabilities are likely to be extremely close to normal traffic. Therefore, misuse detection is not applicable as it will be very difficult to define

a signature. Anomaly detection might detect abnormal activity patterns, but this will be very dependent on the actual deployment case. In any case, since most vulnerabilities are remotely accessible, network-based intrusion detection is the current best tool to detect attacks exploiting these vulnerabilities. Detection of malicious code is only feasible in large devices which include a fully functional operating system (which on the other hand contradicts the protect principle of least functionality PR.PT-3).

5.4 The Respond Function

The Respond Function supports the ability to contain the impact of a potential cybersecurity incident. Examples of outcome Categories within this Function include: Response Planning; Communications; Analysis; Mitigation; and Improvements.

Advisories should provide relevant information for the respond function, as they include a remediation (and/or mitigation) section. However, information in this section is often of uneven quality. The general recommendations provided by centers such as CERT are often too broad to be applicable to a specific sector such as healthcare in a straightforward manner. The recommendations provided by vendors vary greatly. While this function should be present in advisories, this is clearly an avenue for progress.

5.5 Recommendations

Going further, the first recommendation is to study and provide solutions for authentication in healthcare. As the study demonstrates, the main issue is for patients and care staff to obtain easy access to devices and software platforms, even across shifts or in unusual situations. This requirement is accrued by the need for emergency access, that should be included in novel access control frameworks. At the same time, a subset of the devices will not be able to include complex authentication procedures, when they require significant computing, storage, or bandwidth capabilities. Seamless, transparent authentication that offers easy access in all situations while ensuring accountability and post-mortem analysis is a mandatory step forward.

The second recommendation is practical network access control. Many vulnerabilities are remotely exploitable. However, legacy equipment will rely on legacy protocols for many years, and ease of use will also (as with authentication) require simple protocols. Therefore, solutions based on Software Defined Networking are likely to offer the required degree of agility and granularity in network access control that

will be required for healthcare applications. The difficulty is the deployment and management of such technologies, and their integration in current network architectures. In that respect, the work that seems to be starting at the IETF, leveraging protocol sets for cybersecurity, is an interesting avenue to pursue.

The third recommendation is related to data. Healthcare is and will continue to be a data-intensive domain, both for care (e.g. patient monitoring, surgery, diagnosis) and for research to design and evaluate new treatments. Therefore, significant care should be taken that data remains private, accurate and available when needed. The data aspect of vulnerabilities seems undervalued at this stage, and new methods for generating, storing, transporting and using data should be developed that enhance confidentiality, integrity and availability of said data.

6 CONCLUSIONS

There exists a significant volume of information related to cybersecurity in healthcare. Since it is considered a critical infrastructure, there is a lot of information in advisories, in general of better consistence and quality than is available in journalistic sources. This information is reliably structured, up to a point where it is possible to envision further work to analyze this information. The ICS-CERT and NVD databases have formed the basis for creating a dataset of information about the Internet of Medical Things, and medical software, that constitute the healthcare critical infrastructure. Since the requirements to report cybersecurity issues in the critical infrastructure sector have been enacted, there has been a significant number of advisories reported. There are however few vendors involved, so there is a need to check whether this is normal or not. Advisories may regroup several vulnerabilities on several products of the same vendor. Many of the vulnerabilities are rated high or critical, meaning that they have a CVSS score above 7 or above 9.

The analysis demonstrate a take-up of reporting in 2017, associated also with a change of the advisory format. This is a clear response to the worldwide deployment of cybersecurity reporting regulations such as the NIS directive (NIS Directive, 2016) in the EU.

Authentication issues represent the main source of vulnerability present in devices, by far. This seems rather normal in a healthcare environment where access to the device should be given to staff easily, where there is little culture for protection and access control, and where the environment is considered trustworthy. It would probably be useful to reinforce

good programming practices to avoid common errors, and also to ensure that software remains sufficiently simple in devices to be analyzed by existing tools. As expected, the software fault pattern profile for devices is different from the one of software products. These software products generally exhibit data manipulation issues of various origin, that enable the attacker to compromise the software through code injection.

Devices seem to be easily exploitable over the network, with a low complexity. This is coherent with the observation that authentication issues are the most prevalent, because we have experience of default password sharing lists or backdoor information being widely shared for a long period of time. The major impact of vulnerabilities on devices is on confidentiality, as only 10% list no impact and 72% list a high impact. Integrity and availability are slightly less impacted. Also, one has to note that the impact of vulnerabilities on software is generally less and less widespread than the impact of vulnerabilities on devices. Half of the device vulnerabilities list a high impact in the 3 dimensions, while this is the case for only 32% of the pure software vulnerabilities.

The devices have a different patching profile than pure software. On one hand this is not surprising because devices require more effort to patch. On the other hand, medical devices (contrary to the ones found in for example smart homes) are managed by professional, and contrary to other settings (e.g. industry) it is possible to have shorter usage lifecycle that could fit a patching model. Furthermore, critical infrastructure has a requirement to maintain these devices in order to remediate cybersecurity issues. The gap is closing in recent years, but effort is still required to continue in this direction.

The findings were then placed in the context of the NIST Cybersecurity Framework, which provides a standard representation for improving the cybersecurity of critical infrastructure. Out of the five (Identify, Protect, Detect, Respond, Recover) functions of the framework, the first and the last are mostly organizational. The analysis shed only a limited light on the recommendations of the framework, demonstrating mostly that continuous improvement in cybersecurity is also shown in advisories.

ACKNOWLEDGEMENTS

This work was performed while Hervé Debar was a visiting professor with the Center for Trustworthy IoT Infrastructure at Japan Advanced Institute of Science and Technology (JAIST) in Ishikawa, Japan.

REFERENCES

- Annas, G. J. et al. (2003). Hipaa regulations—a new era of medical-record privacy? *New England Journal of Medicine*, 348(15):1486–1490.
- Boeckl, K., Fagan, M., Fisher, W., Lefkowitz, N., Megas, K., Nadeau, E., Piccarreta, B., O'Rourke, D. G., and Scarfone, K. (2019). Considerations for managing internet of things (iot) cybersecurity and privacy risks. *National Institute of Standards and Technology, NISTIR 8228*.
- Coventry, L. and Branley, D. (2018). Cybersecurity in healthcare: A narrative review of trends, threats and ways forward. *Maturitas*, 113:48–52.
- Fu, K. and Blum, J. (2013). Inside risks controlling for cybersecurity risks of medical device software. *Communications of the ACM*, 56(10).
- Kramer, D. B. and Fu, K. (2017). Cybersecurity concerns and medical devices: lessons from a pacemaker advisory. *Jama*, 318(21):2077–2078.
- Kruse, C. S., Frederick, B., Jacobson, T., and Monticone, D. K. (2017). Cybersecurity in healthcare: A systematic review of modern threats and trends. *Technology and Health Care*, 25(1):1–10.
- Marinovic, S., Craven, R., Ma, J., and Dulay, N. (2011). Rumpole: a flexible break-glass access control model. In *Proceedings of the 16th ACM symposium on Access control models and technologies*, pages 73–82. ACM.
- Martin, G., Kinross, J., and Hankin, C. (2017). Effective cybersecurity is fundamental to patient safety.
- Martin, R. A. (2007). Common weakness enumeration. *Mitre Corporation*.
- Mell, P., Scarfone, K., and Romanosky, S. (2006). Common vulnerability scoring system. *IEEE Security & Privacy*, 4(6):85–89.
- NIS Directive (2016). Directive (eu) 2016/1148 of the european parliament and of the council of 6 july 2016 concerning measures for a high common level of security of network and information systems across the union. *OJ L*, 194(19.7):2016.
- Shen, L. (2014). The nist cybersecurity framework: Overview and potential impacts. *Scitech Lawyer*, 10(4):16.
- US Department of Homeland Security (2019). Ics-cert advisories. Online.