Article

Not peer-reviewed version

Assessing Cybersecurity of Internet-Facing Medical IT Systems in Germany & Spain Using OSINT Tools

Pere Tuset-Peiró *, Michael Pilgermann, Josep Pegueroles, Xavier Vilajosana

Posted Date: 18 March 2025

doi: 10.20944/preprints202503.1340.v1

Keywords: medical IT systems; PACS; DICOM; OSINT

 doi:10.20944/preprints202503.1340.v1

Article

Assessing Cybersecurity of Internet-Facing Medical IT Systems in Germany & Spain Using OSINT Tools

Pere Tuset-Peiró [1,*], Michael Pilgermann [2], Josep Pegueroles [3] and Xavier Vilajosana [4]

[1] TecnoCampus-Universitat Pompeu Fabra, Av. Ernest Lluch, 32, 08302 Mataró, Spain
[2] Brandenburg University of Applied Sciences, Magdeburger Str. 50, 14770 Brandenburg an der Havel, Germany; michael.pilgermann@th-brandenburg.de
[3] Universitat Politècnica de Catalunya, C/Jordi Girona, 1-3, 08034 Barcelona, Spain; josep.pegueroles@upc.edu
[4] Universitat Oberta de Catalunya, Rambla del Poblenou 156, 08017 Barcelona, Spain; xvilajosana@uoc.edu
* Correspondence: ptuset@tecnocampus.cat

Abstract: This paper investigates cybersecurity threats in medical IT (Information Technology) systems exposed to the Internet. To that end, we develop a methodology and build a data processing pipeline that allows to gather data from different OSINT (Open Source Intelligence) sources, and processes it to obtain relevant cybersecurity metrics. To validate its operation and usefulness, we apply it to two countries, Germany and Spain, allowing to study the main threats that affect medical IT systems in these countries. Our initial findings reveal that 20% of German hosts and 15% of Spanish hosts tagged as medical devices have at least one CVE (Common Vulnerabilities and Exposures) with a CVSS (Common Vulnerability Scoring System) graded as critical (i.e., value 8 or greater). Moreover, we found that 74% of CVEs found in German hosts are dated from earlier than 2020, whereas for Spanish hosts the percentage is 60%. This indicates that medical IT systems exposed to the Internet are seldom updated, which further increases their exposure to cyberthreats. Based on these initial findings, we finish the paper providing some insights on how to improve cybersecurity of these systems.

Keywords: medical IT systems; PACS; DICOM; OSINT

1. Introduction

In the early 1980s, medical equipment manufacturers pioneered interoperability by developing standards, such as ACR/NEMA 300, to facilitate data exchange between CT (Computed Tomography) and MRI (Magnetic Resonance Imaging) systems. The advent of the Internet in the early 1990s further enhanced communication and management of medical devices, with medical image protocols adopting the TCP/IP stack. At that time DICOM (Digital Imaging and Communications in Medicine) emerged as a prominent protocol enabling storage, retrieval, and printing of medical images across various devices [1]. To further ease the integration of DICOM devices with the Internet, PACS (Picture Archiving and Communication Systems) systems were developed in the late 1990s , which integrate a DICOM interface to retrieve data from medical devices and provide a common HTTP (Hyper-Text Transfer Protocol) interface to expose data to medical staff and patients through a web browser [2].

Such approach has clearly increased the efficiency in patient care, but it has also introduced significant cybersecurity risks, specially when considering that integrating security mechanisms to DICOM and HTTP was not a priority in the early days. This has set the stage for evolving cybersecurity risks in the medical domain, specially when considering that these medical IT systems are now being exposed to the Internet. For example, in 2016 researchers discovered 2774 DICOM servers exposed to the Internet worldwide [3]. In 2019 researchers from Greenbone Networks [4] showed that many PACS servers worldwide are accessible via the Internet without any security measures. In 2023 Sina Yazdanmehr from aplite GmbH also showed in a BlackHat conference that DICOM servers and HTTP servers used in medical IT systems are largely exposed to the Internet, which can lead to patient data leaks. Last but not least, in 2024 Censys [4] identified over 14000 unique IP addresses exposing

 doi:10.20944/preprints202503.1340.v1

medical-related services, including 5100 DICOM and 4000 EMR/EHR (Electronic Medical Records / Electronic Health Records) services. Hence, there is an urgent need to comprehensively assess the cybersecurity status of Internet-exposed medical IT systems around the world.

Given these preliminary findings, this study aims to propose a flexible and open data acquisition and processing pipeline that enables the collection and analysis of information from healthcare systems connected to the Internet. The proposed methodology will allow for the investigation of vulnerabilities in medical IT systems connected to the Internet based on OSINT (Open Source Intelligence) tools. We decide to use OSINT tools, such as Censys and Shodan, since its usefulness has already been demonstrated to study vulnerabilities of Internet-facing systems, such as power grid systems [5]. Using the pipeline we focus this initial study on German and Spanish medical IT systems, allowing to validate its correct operation and to gather evidence about their status. However, the methodology and the tools employed in this study can serve as a foundation for wider studies assessing global medical IT cybersecurity, which we plan to do in the future. Last but not least, it is important to disclose that during the development of this project no active hacking has been conducted. Moreover, all the results that have been found are in the process of being reported to the stakeholders, process that will be completed before the publication of this study.

The remaining of the paper is organized as follows. Section 2 presents an overview of cybersecurity threats faced by medical IT systems. Section 3 details the methodology that we have followed and the tools that we have developed to conduct this study. Section 4 presents and discusses the results that have been obtained as part of the study. Finally, Section 5 summarizes our findings and outlines future work.

2. Cybersecurity in the Medical Domain

There is limited generic material available on how hospitals look with respect to their business processes and their technological implementation.

In general, (networked) systems in a hospital are grouped along the following domains: IT (Information Technology), MT (Medical Technology), and utility systems. In turn, the systems in the IT domain can further be divided into two categories: general administrative IT (PCs and servers that support administrative tasks and communication) and specialized hospital IT systems, including HIS (Hospital Information System), RIS (Radiology Information System), LIS (Laboratory Information System), and PACS (Picture Archiving and Communication Systems). MT in hospitals includes a wide range of critical devices used for diagnosing, monitoring, and treating patients. These include ventilators, infusion pumps, patient monitoring and imaging systems such as MRI (Magnetic Resonance Imaging) and CT (Computed Tomography) scanners. Due to their growing network connectivity, these devices are increasingly referred to as Medical Cyber-Physical Systems (MCPS) in the literature [6].

The heterogeneity of systems across the IT and MT domains in hospitals creates significant complexity. Furthermore, the interconnection of these devices within and between these domains (including the Internet) offers numerous operational advantages, but also introduces new security vulnerabilities that must be carefully managed to ensure the safety of both patients and staff. Hence, detecting and patching vulnerabilities throughout the medical equipment in a hospital has become a top priority for the staff, but the teams managing such processes are often non-existing or overloaded, thus delaying such procedures.

2.1. Hospitals IT Cybersecurity Status

Sources like Check Point Research publications [7] report that the healthcare sector is now among the three most attacked ones; and ENISA (the European Union Agency for Cybersecurity) claims that hospitals are particularly affected in terms of cybersecurity, as RansomWare remains the predominant threat, driven by the prospect of financial gain [8]. Following reports of the Health Sector Cybersecurity Coordination Center, the interest of attackers has increased to the point that the that the top ten ransomware groups actively target the US health and public health sector [9].

 doi:10.20944/preprints202503.1340.v1

2.2. Vulnerabilities and Their Management

Regarding to [10], 33.524 vulnerabilities were recorded in the NIST NVD between July 2023 and June 2024, which makes up for an increase of 35 % compared to the reporting period one year before (24.690 vulnerabilities). For organisations it is challenging to identify the most relevant vulnerabilities and prioritize countermeasures. To address these issues, VM (Vulnerability Management) was developed [11].

3. Methodology

In this section, we present the proposed methodology, along with the existing tools selected and those developed specifically for this work. In each corresponding subsection the detailed steps of each phase is described in detail.

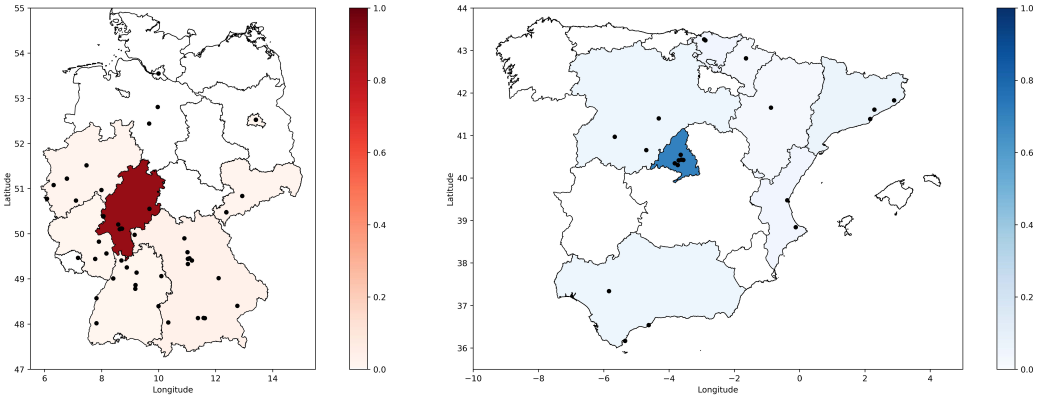

(a) Germany map (b) Spain map

Figure 1. Pipeline.

3.1. Scanning the Internet for Medical IT Systems

Scanning the Internet to discover services running on hosts has been one of the main ways to uncover vulnerabilities in different environments, albeit a slow one. Indeed, scanning the whole range of IPv4 addresses is an endeavour that can take various weeks if done sequentially given the large number of hosts (i.e., $2^{32} = 4294967296$), ports (i.e., $2^{16} = 65536$) and protocols (i.e., TCP and UDP) to explore. The process can be speed up by splitting the work among various servers and running it in parallel, but it is still a slow process given the magnitude of the search space. In addition, while scanning hosts to detect services is not considered a crime in the United States and Europe, it can still cause legal problems when such procedure is not explicitly authorized by the administrators. Moreover, active scanning is considered an attack by many SOCs and cause alarms when done massively.

Since we decided to adhere to a strict no-active hacking for the development of the project, we use a methodology based on OSINT, which stands for Open Source Intelligence and, according to its pioneer Robert David Steel, it is "*the art and science of creating ethical, evidence-based decision support using only open sources and methods—legal and ethical in every respect*" [7]. Hence, we decided to use search engines that periodically scan the whole range of IPv4 addresses and store the services running in each host, as well as related metadata (i.e., geographical information, DNS information, etc.), allowing to speed up the process and save computing and network resources. Based on the information already stored in their database, these services allow users to perform queries and filter the results to search for specific information. For example, it is possible to query the database to discover all the ports open on a specific IP addresses, or to query for any IP address that is running a certain service. Moreover, since the information stored in the database is augmented with context information (i.e., IP geo-location), it is also possible to query by the country where the hosts are running, among others.

The OSINT tools that we aim to integrate in our processing pipeline should meet the following requirements: 1) Well-established within the community, providing reliable results, 2) Provide basic

 doi:10.20944/preprints202503.1340.v1

information (i.e., ports, services, etc.) and additional metadata (CVEs, geoposition, etc.), 3) Provide frequent updates about the hosts and the services they are running, and, 4) Provide an API to automate the process of making queries and retrieving results.

As of today, there are two main OSINT tools that meet the described criteria: Shodan (https://www.shodan.io/) and Censys (https://censys.com/). The former was founded by John Matherly in 2009, whereas the latter was founded by Zakir Durameric in 2017, one of the developers of ZMap [8], a collection of open source tools to perform large-scale studies of Internet hosts and services. Shodan provides more detailed information for each host and service, like a list of CVEs, which is very useful to report found vulnerabilities . In contrast, Censys has a better query and filter system, allowing to detect hosts and services with more precision.

3.2. Evaluating OSINT Tools Through a Hackathon

In order to explore the problem at hand and evaluate the OSINT tools, we organized a hackathon as an activity of the CARISMATICA project, led by UPC (Universitat Politècnica de Catalunya) and funded by INCIBE (Instituto Nacional de Ciberseguridad). The hackathon took place at TecnoCampus Mataró in December 2024, with the participation of 18 students from various universities in Germany and Spain. The participants had different backgrounds, including Computer Science and Electrical Engineering, and different academic levels, from Bachelor's to Master's degree.

When preparing the hackathon we devised two approaches to kick off the data collection and processing pipeline:

- *A-priori methodology*: The process starts with a preparation phase, in which IP ranges linked to medical organizations are identified using additional OSINT sources (i.e., an atlas of medical institutions in each country), so that the follow-up steps of the pipeline only work on a subset of the complete public IP address range.
- *Zero-knowledge methodology*: The process starts with a Censys request and solely relies on the tagging mechanisms provided by Censys to find IP addresses linked to medical IT devices. Afterwards, the information is enriched with other OSINT tools, such as Shodan, to obtained a detailed overview of each IP address.

During the hackathon, students were mixed in 4 groups and worked in 2-hours sprints over the course of 4 days, with the only rule that no active hacking was allowed. In the first sprint, we provided some general ideas on the objective and the methodology, as well as access to the OSINT tools described earlier, and allowed them to explore the methodologies and the APIs of each tool. After the first sprint each group selected a methodology, which they would use for the following sprints, and started developing code to implement the data collection and processing pipeline. At the end of each sprint each group of students provided a short update of their progress and any relevant findings, and at the end of the hackathon each group presented their results.

Overall, the results of the hackathon were very positive, both in terms of student engagement and practical outcomes. First and foremost, we could witness how students started from a very basic idea of the goal and almost no knowledge of the tools, to being able to develop code that allowed to get some initial results. It is needless to say that the moment that students got their first results, the energy and the motivation in the room boomed. Second, it allowed to validate our ideas on the data collection and processing pipeline, and to validate the potential of the project. Hence, we believe that hosting a hackathon was a very good experience for everyone involved, and we plan to repeat it next year.

3.3. Developing a Data Acquisition and Processing Pipeline

As we learned through the hackathon, collecting and processing data from OSINT tools involves a huge amount of data from different sources. To simplify data handling and make results reproducible, we have developed a set of processing pipelines using the Python programming language. Several processing steps have to be combined, usually by piping information between several Python scripts or persist output from a processing step into SQLite databases, CSV files, or JSON files. The processing

 doi:10.20944/preprints202503.1340.v1

pipelines have in common that they are meant to work on a set of IP addresses, which are linked to medical IT devices, and run (semi-automatically) several OSINT procedures, allowing to identify and report on relevant security information for each host.

Based on the experience acquired through the hackathon, we decided to use the *zero knowledge* methodology. Hence, we start by using Censys to collect the IP addresses of hosts in each country that are tagged as '`medical-device`'. To that end, we developed a Python script that uses the Censys API search feature to run the query '`location.country: 'country' and labels='medical-device''`, where `country` is a variable indicating the country (i.e., Germany or Spain). The results returned by Censys are in the form of a JSON structure with the IP addresses of devices that meet the search criteria. After that, we use the Censys API view feature to retrieve all the information for each host. For each host, the Censys API returns a JSON structure that contains information about the host in the form of a dictionary. The keys of the dictionary point to other data structures that can be other dictionaries or lists. At the first level of the structure, the dictionary includes keys such as `ip`, `services`, `location`, `whois` and `dns`, among others. Then, within the `services` key we find a list of dictionaries, each containing further information about each services, such as, the name of the service (`service_name`) and the port in which the service is running (`port`), among others. All the information of each host retrieved in each query is stored in a file that is named after the IP address of the host and with the `.censys` extension. The files are stored under a directory named after the country, i.e., '`Spain/IP_address.censys`'.

Once all the hosts for each country have been retrieved using Censys, another Python script uses the Shodan API host feature to query information about that host. The information provided by Shodan is structured in a similar way (i.e., dictionary containing dictionaries or lists), but provides different information. For example, the Shodan query returns a first level dictionary with a summary of the collected information (i.e., IP address, ISP, geolocation, ASN, etc.) and another dictionary with all the raw data. Again, the information retrieved from Shodan is stored in a file named after the IP and with the `.shodan` extension under the country directory, i.e., '`Spain/IP_address.shodan`'. During the process we have found that in some cases Shodan does not have information for a given host provided by Censys, in which case the files are not created. This is accounted for when further processing the data files of each country to calculate statistics and plot the results.

Once all the data for each country has been retrieved from Censys and Shodan, another Python script is responsible to analyse the services and ports running in each host and to grab screenshots of the HTTP services using the Selenium library. The screenshots of each host are stored in a PNG file, allowing to manually analyse the obtained results for each host/port and determine which medical IT services are running on each.

Finally, another set of Python scripts in the data processing pipeline are responsible to analyse the data, compute statistics and create the plots, allowing us to analyse the potential impact of vulnerabilities, compare the state of cybersecurity in each country, and to draw conclusions. In the process of analysing the potential impact of vulnerabilities, it is important to mention that we use the CVE (Common Vulnerabilities and Exposure) system and the CVSS (Common Vulnerability Scoring System) score. On the one hand, the CVE is a unique number that uniquely identifies each discovered vulnerability, either in software or hardware. On the other hand, the CVSS is a number between 0 and 10 that is assigned to each CVE depending on its severity level, ranging from low (0.0-3.9) to critical (9.0 to 10.0). To find information about the CVE and the CVSS we use the NVD (National Vulnerability Database), which receives its information directly from MITRE and is considered to be a primary source [9,10].

4. Results

This section presents and analyses the results that have been obtained using the methodology described earlier in the paper.

 doi:10.20944/preprints202503.1340.v1

4.1. Dataset Overview

The procedure to gather information from German and Spanish hosts tagged as medical devices was conducted the third week of February 2025. It is important to mention that the process of gathering the data is slow, since both Censys and Shodan impose a limit of around one query per second and the lists of host can be large (i.e., over 1000 per country).

Using the methodology described earlier, Censys returned a total of 54 hosts for Spain, and a total of 1449 hosts for Germany. In turn, Shodan returned 49 hosts out of the 54 hosts (90.7%) for Spain, whereas it returned 1426 hosts of the 1449 hosts (98.4%) for Germany.

Based on the obtained results, we start our dataset overview by checking on information freshness of each service by comparing the `last_updated_at` tag on Censys and the `last_update` tag on Shodan. For Spain the results show that 75% of the time information provided by Censys is more up-to-date than Shodan. In contrast, for Germany the results show that 75% of the time Shodan provides more up-to-date information than Censys.

We also check the size of the information (i.e., JSON structure) returned by each service for each host. The results for Spain show that on average Censys returns two times (2x) more information for each host than Shodan, whereas for Germany Censys returns six times (6x) more information for each host than Shodan. However, for some hosts we have found that Shodan provides more detailed information, such as a list of the CVEs affecting the hosts.

Finally, we also show were hosts are located in each country, as depicted in Figure 2. The black dots in the maps show the (approximate) location of each hosts, whereas the heat-map shows the percentage of hosts found in each region of the country with respect to the total number of hosts detected in the country. As it can be observed, the regions of Hessen in Germany and Madrid in Spain are the ones with a higher percentage of vulnerable hosts found, with a percentage of 90% and 70%, respectively. Such concentration can be explained by the fact that most services are hosted by cloud providers, such as Amazon (i.e., the `3.64.0.0/12` IP range is announced from Hessen), or Internet Service Providers, such as Telefonica (i.e., in Madrid).

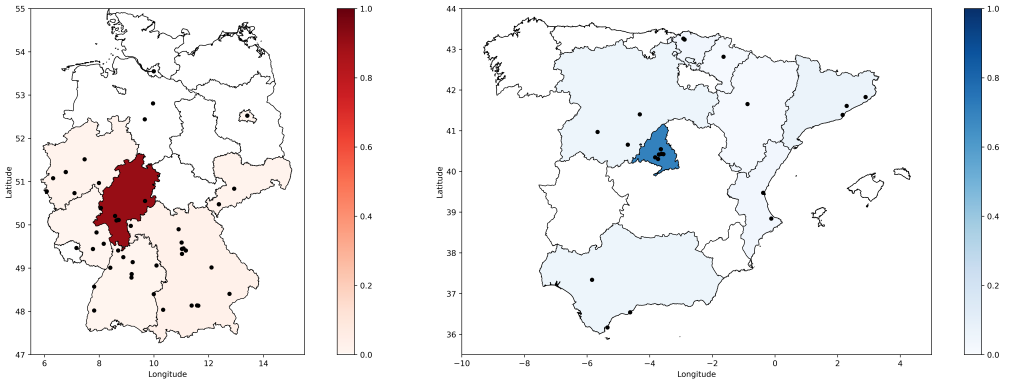

(**a**) Germany map (**b**) Spain map

Figure 2. Heat-map of Germany and Spain showing where hosts tagged as medical devices have been geo-located. Most hosts have their origin in the regions of Hessen (Germany) and Madrid (Spain), where cloud datacenters are located.

4.2. Dataset Cleaning

To analyse the potential cybersecurity impact on medical devices in Germany and Spain, we start by focusing on the ports that are open for every host. Analysing the dataset we notice that in Germany the average number of open ports per host is 52.86 ports, with a standard deviation of 56.88, and the host with the most number of open ports is 582. In contrast, in Spain the average number of open ports per host is 7.94, with a standard deviation of 7.35, and the host with the most number of open ports is 37.

 doi:10.20944/preprints202503.1340.v1

Having hosts with such large number of open ports makes us suspicious that a fraction of hosts may be honeypots. Hence, we start by plotting the percentage of hosts with a given number of open ports for Germany (red) and Spain (blue). The results depicted in Figure 3 confirm our suspicion. As it can be observed, in Germany the distribution shows two clear groups of hosts: a group with a number of open ports between 1 and 40, and another group with more than 40 open ports. In contrast, in Spain all hosts have less than 30 open ports except one. However, only relying on the number of open ports may lead to some hosts being flagged as false positives (i.e., regular hosts flagged as honeypots) or false negatives (i.e., honeypots flagged as regular hosts).

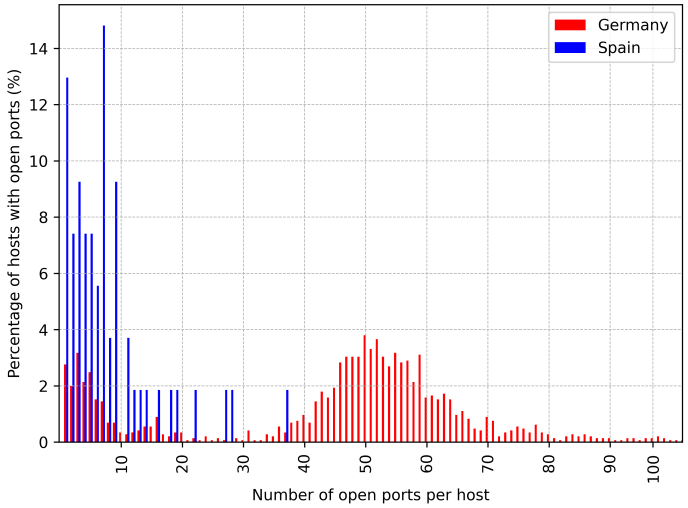

Figure 3. Percentage of hosts tagged as medical devices with a given number of open ports per host in Germany and Spain.

To further confirm our hypothesis, we analyse the hosts in Germany to determine where the IPs of the hosts are being originated from. As depicted in Figure 4, we observe that most hosts that have more than 30 open ports are originated from ASN 16509, which belongs to AWS (Amazon Web Services).

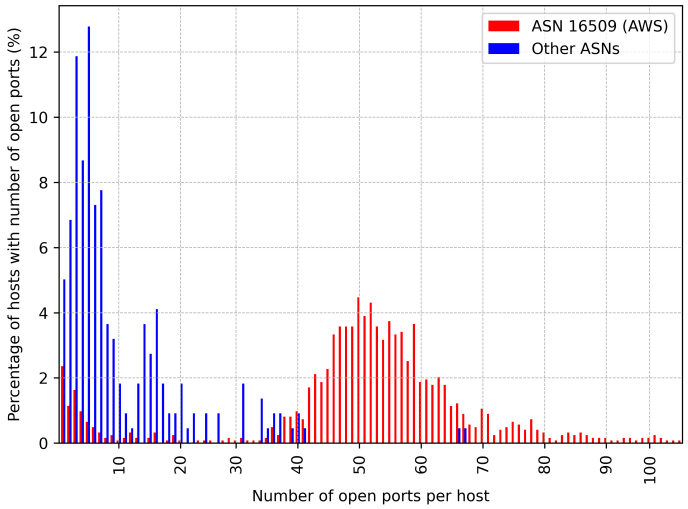

Figure 4. Percentage of German hosts tagged as medical devices with a given number of open ports per host classified by the ASN where the host IP address is originated from.

In addition to that, we also check the values returned from the reverse DNS query on the IP address. Using such approach, hosts that are not honeypots should only have a handful of DNS records to allow humans access the services that are hosted. In contrast, hosts that are honeypots should have

 doi:10.20944/preprints202503.1340.v1

none or only one DNS record, since humans are not expected to access the services that are hosted anyway. Moreover, as we have seen earlier, hosts that are honeypots will most probably be announced from a cloud service provider, such as AWS. As depicted in Figure 5, hosts in Germany with over forty open ports and zero or one DNS record are mostly announced from AWS, confirming our hypothesis.

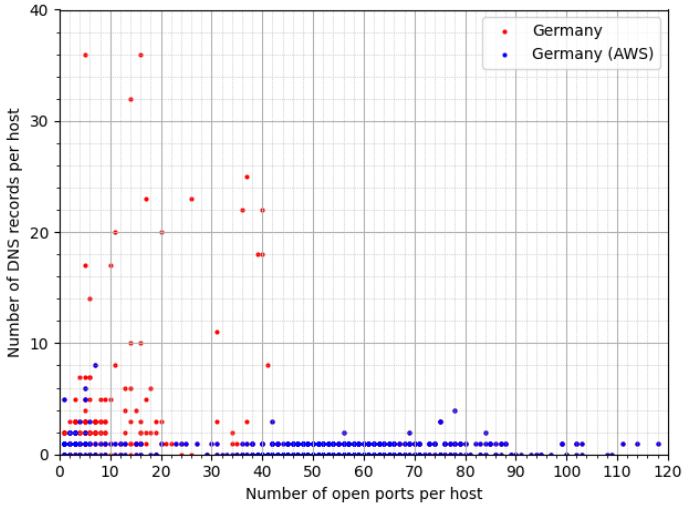

Figure 5. Correlation between the number of open ports per host and the number of DNS records per host for Germany.

Hence, seeing the strong correlation between the three approaches, we have decided to remove hosts with more than 30 open ports from the remaining of our analysis, since they can have a huge impact on the results and the conclusions.

4.3. Dataset Analysis

After removing the hosts flagged as potential honeypots, we now focus on the services, as determined by the protocol that is executing (and not the port where the socket is listening), that are run by hosts in Germany (red) and Spain (blue), as depicted in Figure 6. As it can be observed the most used protocol in hosts tagged as `medical-devices` by Censys is HTTP, with 53% for Germany and 42% for Spain. The remaining results for Germany and Spain are similar, with hosts running SSH, DICOM, FRP, IMAP, MYSQL, POP3 and SMTP, among others. In addition, there are also a 12.8% of ports that are classified as unknown.

Having HTTP services exposed to the Internet is normal as long as such services are intended for end-users accessing it, and as long as the services are implemented following appropriate coding guidelines. What is not normal is having so many services running on such hosts. For example, we observe some hosts exposing database or email services.

 doi:10.20944/preprints202503.1340.v1

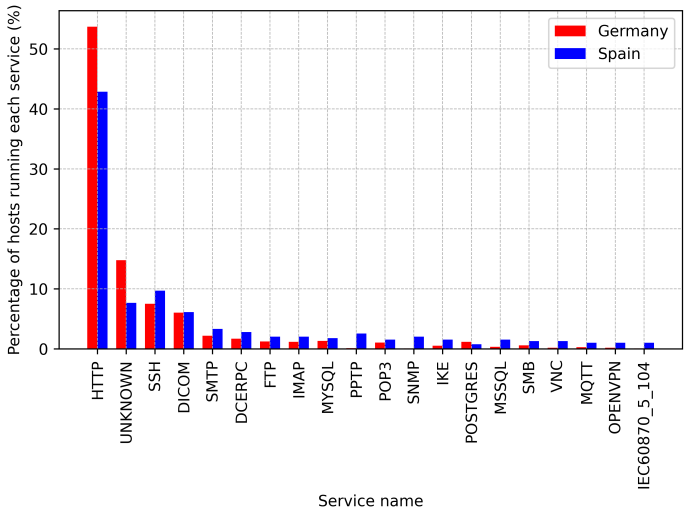

Figure 6. Percentage of most common services run by hosts tagged as medical devices in Germany and Spain.

Next, we focus on analysing the operating system used by hosts, since this is another potential threat surface, specially for hosts running older operating systems that may not be patched. As displayed in Figure 7, in Germany we observe that 35% of the hosts use Linux, whereas Windows only represents 4% of the hosts. In contrast, in Spain Linux represents 28% of the hosts, whereas Windows represents 15% of the hosts. Unfortunately, for percentage larger than 50% of hosts from German and Spain the operating system could not be correctly detected. One possible explanation is that some hosts are running behind a firewall implementing destination NAT, so properly identifying the operating system of each host is not an easy task since different ports can refer to different hosts running different operating systems. Moreover, it has not been possible to collect additional information regarding the operating system of each host (i.e., version, kernel, etc.), since it is not provided by either Censys or Shodan.

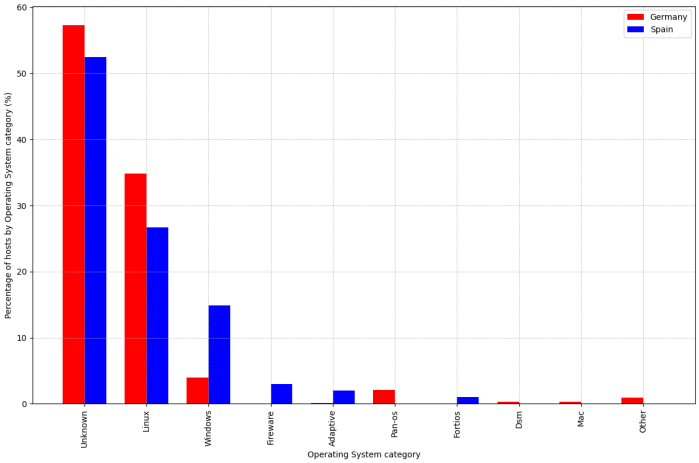

Figure 7. Percentage of most common operating systems run by hosts tagged as medical devices in Germany and Spain.

We now focus on the vulnerabilities found in the services running on each host. In particular, we analyse the CVEs that affect each host/service, and the CVSS score that determines the severity of the vulnerability. This is an important aspect for hosts exposed to the Internet, since an exploitable vulnerability (i.e., buffer overflows, SQL injection, etc.) can render the authentication of the server useless. We do so by looking at the '`vulns`' tag provided by Shodan for each service.

Figure 8 shows the percentage of German (red) and Spanish (blue) hosts that belong to each CVSS category, with values ranging from 0 to 10 depending on the severity of the vulnerability. As it can be

 doi:10.20944/preprints202503.1340.v1

observed, more than 20% of German hosts and more than 15% of Spanish hosts are affected by at least one vulnerability with a CVSS value between 9.0 and 10.0, indicating a critical vulnerability that is easy for attackers to exploit and has a large compromise potential of the target system. Moreover, more than 40% of German hosts and more than 30% of Spanish hosts are affected by at least one vulnerability with a CVSS between 7.0 and 8.9 representing a vulnerability that is harder to exploit but which leads to a high compromise potential on the target.

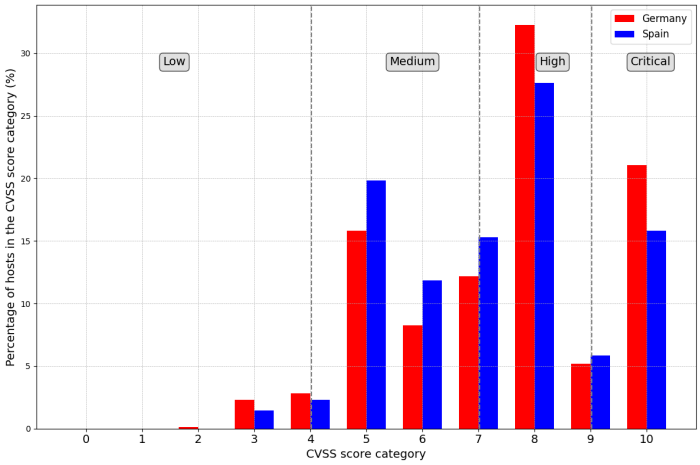

Figure 8. Percentage of hosts tagged as medical devices in German and Spain with CVEs that belong to each CVSS category (i.e., low, medium, high, and critical).

In addition, Figure 9 shows the percentage of German (red) and Spanish (blue) hosts with the top 30 CVE vulnerabilities found and their related CVSS score. As it can be observed, most of the top 30 CVE vulnerabilities found have CVSS values lower than 8, indicating that they are either difficult to exploit or that they lead to a small compromise of the target host/service. However, we also observe various CVE vulnerabilities with CVSS scores larger than 8, indicating that they are easy to exploit or that they can lead to a large compromise of the target system. These CVEs and their associate CVSS score are the following: CVE-2024-38476 (CVSS=9.8), CVE-2024-38474 (CVSS=9.8), CVE-2021-44790 (CVSS=9.8), CVE-2022-31813 (CVSS=9.8), CVE-2021-39275 (CVSS=9.8), CVE-2022-22720 (CVSS=9.8), CVE-2022-22721 (CVSS=9.1), CVE-2021-40438 (CVSS=9.0). It is worth noting that most of the CVEs found affect the Apache HTTP Server leading to DOS (Denial of Service) or RCE (Remote Code Execution).

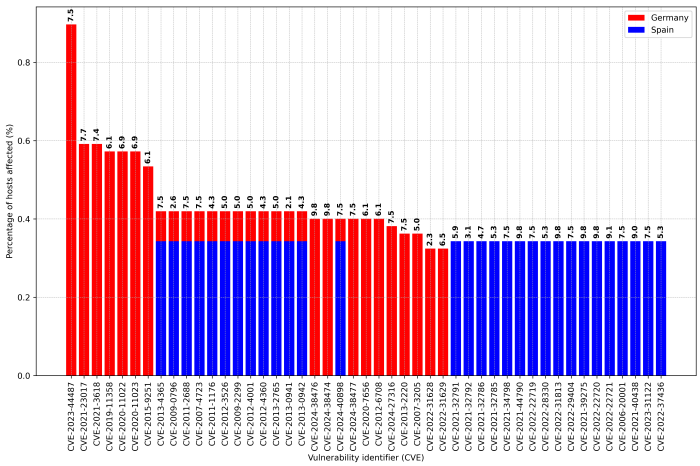

Figure 9. Percentage hosts tagged as medical devices in Germany and Spain with the top 30 vulnerabilities found and their related CVSS score.

 doi:10.20944/preprints202503.1340.v1

Figure 10 shows the distribution of CVEs found in German (red) and Spanish (blue) hosts classified by the year that they were discovered and classified into the CVE system. As it can be observed, there is a large percentage of CVEs that found in these hosts that have a date In particular, the percentage of CVEs found in German hosts that date from earlier than 2020 is 74%, whereas for Spain that number is about 60%.

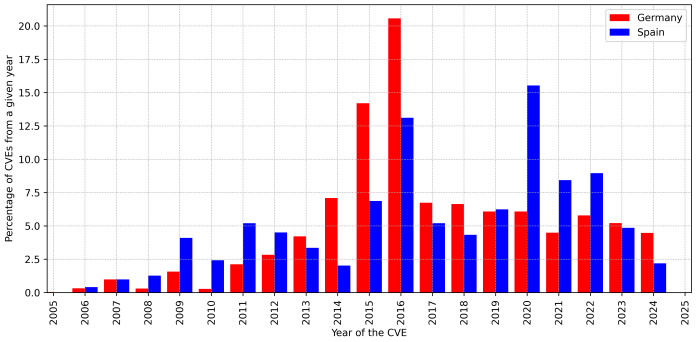

Figure 10. Percentage of CVEs hosts tagged as medical devices in Germany and Spain classified according to the year where they were discovered.

To finalize our analysis we study the screenshots that we have obtained for each port that runs an HTTP service on each host. Using such approach, we have detected several installations with default parameters, accessible without credentials, and meant for demonstration purposes. We provide one example of each category that we have found in Germany. First, the host `62.171.155.131` runs an Orthanc installation with a insecure setup warning. Second, the host `82.165.144.45` runs an Orthanc ICQ Cloud software that is accessible without credentials. Finally, the host `212.132.115.76` is pointed by the DNS record `demo.radio.kapsiki.net`, indicating that this is not running a real medical IT service.

Analysing the remaining results we have also found several hosts that are worth describing. As a summary of the results, we only focus on three hosts running in Spain that have been tagged as medical devices. The idea is to provide a general overview of the current state of cybersecurity in hosts tagged as medical devices, allowing to discuss the threats and the potential solutions in the following subsection.

The first host `2.136.109.191` belongs to an image diagnostics centre located in Catalunya, as obtained using a reverse DNS query (i.e., `pacs.imaginebarcelona.com` domain) and the IP geo-position. Analysing the Censys results shows that the host has 4 ports open: 104 (DICOM), 161 (SNMP), 9090 (HTTP), 9091 (HTTP). In addition, analysing the Shodan results for the host shows that there are no known CVEs affecting it. Having a low number of services exposed to the Internet and no known CVEs, which is a good sign of cybersecurity hygiene. However, we are unsure why these ports are even required to be exposed to the Internet, specially DICOM and SNMP which are not meant for end-users accessing a service. Analysing the HTTP services on running ports TCP/9090 and TCP/9091 shows the landing page of a DAIPACS, a PACS software.

The second host `217.127.207.41` belongs to a medical centre located in the Castilla y Leon, as obtained using a reverse DNS query and the IP geo-position. The host has more than 10 ports publicly exposed, and most of these ports are accessible using HTTP. Through a quick analysis of the data provided by Censys and Shodan we determine that these ports expose services such an SCADA server (Tridium Niagara), a DICOM viewer (OHIF) and a database (MySQL). In addition, one of the HTTP services exposes the landing page of a firewall (Sophos UTM), indicating that most of these services run behind a NAT implementing redirection to internal IP addresses based on the destination port. Besides the large number of ports, the most relevant aspect it that the host has been flagged with more than 21 CVEs, with CVSS ranging from low to critical. In particular, the hosts contains the following CVEs (and they related CVSS) that are specially relevant: CVE-2017-8923 (CVSS=9.8), CVE-2017-9120 (CVSS=9.8), CVE-2021-21708 (CVSS=8.2), CVE-2022-31625 (CVSS=8.1), and CVE-2022-37454 (CVSS=9.8). Again,

 doi:10.20944/preprints202503.1340.v1

the vulnerabilities found in the host have a high CVSS score, indicating a high risk potential. Moreover, the CVEs mostly affect the Apache HTTP server or related libraries (i.e., PHP), and can lead to DOS and RCE.

The final interesting host (`83.48.110.205`) that we found in Spain belongs to a medical centre located in the Madrid region, as validated using a reverse DNS query (i.e., `traumasport.com` domain) and the IP geo-position. According to both Censys and Shodan the host runs the Microsoft Windows Server 2012 operating system and only has 6 ports open. Two ports (TCP/80 and TCP/443) run HTTP/HTTPS to host the domain website. In addition, the host also has ports TCP/137 (NetBIOS) and ports TCP/5901 and TCP/5915 (VNC) open to the public. Finally, the host also has port TCP/8999 open, which is linked to the RemotEyeLite Server v2.2.1. Despite the short number of ports, it is already concerning that NetBIOS and VNC are running on ports exposed to the Internet. However, the most eye-catching part is that, despite the short number of ports, the host is vulnerable to more than 200 CVEs. We will not describe each CVE in detail as we have done with other hosts due to the lack of space, but it is enough to say that there are several CVEs dated from 2012 (i.e., CVE-2012-4360 and CVE-2012-4001), 2011 (i.e., CVE-2011-2688 and CVE-2011-4718) and earlier, all of them with CVSS values ranging from low to high.

4.4. Discussion

In the previous subsection we have provided a general overview of the results obtained when querying Censys and Shodan for hosts in Germany and Spain tagged as medical devices, as well as a deeper analysis of three specific hosts located in Spain to showcase the most common vulnerabilities found. While we do not provide an in-deep analysis of hosts tagged as medical devices running in Germany due to lack of space, a quick analysis revealed that the results obtained show a similar situation to the one in Spain. Hence, in this subsection we discuss the cybersecurity outlook of medical IT devices in Germany and Spain, and provide some general guidelines to improve the situation.

4.4.1. Running Unrelated Services in One Host

We have observed hosts tagged as medical IT devices that have multiple ports exposed and run services that are completely unrelated to their category, such as remote control (i.e., VNC) or video recording (i.e. DVR) software, among others. At this point, we should investigate further to determine if the IP address of the host uses a NAT service to redirecting connections to internal IP addresses based on the destination port. However, if it is really a system on which medical IT services are running with other unrelated services, this is a clear violation of good practice for secure IT operations.

4.4.2. Running Services that Are not Patched

We have found hosts that are running services with known CVEs that are quite old (i.e., 10+ years) and have a high CVSS scores (i.e., 8+). Specifically, most of the vulnerabilities that we have detected affect the HTTP server handling the requests (i.e., Apache HTTP server) or related libraries (i.e., PHP) used by the applications. For example, checking the CVEs information we discovered that that the version of Apache that most hosts are running is at least 3+ years old. Hence, while the application itself could be developed according to modern security standards, a bug in the HTTP server or the related libraries may provide a malicious user a mechanism to infiltrate the server and get access to user data.

4.4.3. HTTP Landing Pages for Medical IT Services

Most medical IT services that are exposed to the Internet use the HTTP protocol to allow users connect remotely using a web browser. While most medical IT services found have a login page that requires the user to introduce a username/password to access the service, it is unknown if there are other mechanisms to protect the service behind it. For instance, there could be a default username and password if the installation uses the default configuration, as we have found in some instances through a message in the login page. Also, even if properly configured, there may not be a maximum limit

 doi:10.20944/preprints202503.1340.v1

of login retries or a two-factor authentication mechanism. In fact, judging by the appearance of the user interface, it seems plausible that these limitations do not exist, which is unacceptable considering modern IT security practices.

4.4.4. Re-Using DNS Domains for Different Purposes

We have observed that some medical centres reuse the domain that they use for their main website (i.e., `domain.com`) to point to hosts that run medical IT services by using a subdomain (i.e., `pacs.domain.com`). While this is correct from a technical perspective, reusing the main domain provides a potential attacker with additional information about the owner of the medical services (i.e., where it is located, who are the employees, etc.), which can be exploited to perform other types of social engineering attacks (i.e., phishing). Hence, such practices should be avoided whenever possible.

4.4.5. Improving Cybersecurity of Medical IT Services

As we have just described, there are many threats potentially affecting medical IT services exposed to the Internet.

Hence, to improve cybersecurity of medical IT services we believe that they should be separated from general IT services, and the former should not be accessible from the Internet when possible. Instead, a VPN (Virtual Private Network) should be used to ensure that authorized users can connect to it. The VPN can be a client-to-site VPN in case it is an end-user (i.e., a doctor) that needs to access the service, or a site-to-site VPN in case two medical IT systems have to connect to each other (i.e., a CT or MRI and a PACS). Currently, there are many open-source projects (i.e., WireGuard) that provide simple and affordable ways to deploy such tunnels, so not using them is gross misconduct in terms of cybersecurity hygiene. Moreover, in case of medical IT services that cannot be hidden behind a VPN, we would recommend to: 1) separate the services using a firewall with NAT, 2) separate the domains used for general IT services (i.e., website) from medical IT services (i.e., PACS server). Finally, the operating system and the software running the medical IT services should be frequently patched, and each service should be appropriately hardened, with a limitation on the maximum number of logins or a double-factor authentication to limit the risks of attacks.

Last but not least, we want to stress that the methods used throughout the paper did not involve any active hacking. Moreover, the OSINT tools that we have used are accessible to the general public, only requiring registration and not involving any subscription fees when used for a limited number of hosts. Hence, system/network administrators responsible for medical IT systems should use these tools to gather information about potential security holes caused by lack of updates or misconfiguration, anticipating potential threats before they occur.

5. Conclusions and Future Work

This study has investigated the cybersecurity status of Internet-exposed medical IT systems in Germany and Spain using a processing pipeline that gathers and analyses data from OSINT tools. We found that 20% of German hosts and 15% of Spanish hosts tagged as medical devices have at least one critical vulnerability, with CVSS equal or greater to 8. We also discovered that vulnerabilities found in these systems are old, with 74% of CVEs in German systems and 60% of CVEs in Spain dating from before 2020. Hence, to improve cybersecurity, we recommend applying common IT practices to medical IT systems, including: 1) regular updating of the operating system and the services, 2) implementing strict access control measures, such as using VPNs for remote access, 3) separating general IT services from medical IT systems to minimize attack surface, and, 4) monitoring OSINT data to proactively identify potential vulnerabilities.

Based on our findings, in the future we also plan to: 1) investigate specific vulnerabilities found in this study, providing detailed mitigation strategies and exploitation prevention techniques, 2) monitor the evolution of CVEs identified in our study over time to assess the effectiveness of remediation actions taken by administrators, and, 3) conduct a worldwide study using the developed methodology and tools to assess the cybersecurity status of Internet-facing medical IT systems across various

 doi:10.20944/preprints202503.1340.v1

countries. In addition, we believe it would be valuable to foster the collaboration with with medical IT vendors, healthcare providers, and regulatory bodies, to integrate our findings into their cybersecurity guidelines.

Acknowledgments: This project is carried out within the framework of the Recovery, Transformation, and Resilience Plan, funded by the EU Next Generation funds, under the auspices of the INCIBE cybersecurity chair named CARISMATICA. The authors of the project would also like to acknowledge the collaboration of Censys in the development of the project by providing unrestricted access to their tools.

References

1. DICOM (Digital Imaging and COmmunications in Medicine). https://www.dicomstandard.org. Last accessed: 2025-02-17.
2. Yan, L. DICOM Standard and Its Application in PACS System. *1*. https://doi.org/10.24294/mipt.v1i1.221.
3. Stites, M.; Pianykh, O.S. How Secure Is Your Radiology Department? Mapping Digital Radiology Adoption and Security Worldwide. *American Journal of Roentgenology* **2016**, *206*, 797–804, [https://doi.org/10.2214/AJR.15.15283]. PMID: 26934387, https://doi.org/10.2214/AJR.15.15283.
4. The Global State of Internet of Healthcare Things (IoHT) Exposures on Public-Facing Networks. https://censys.com/state-of-internet-of-healthcare-things/, 2024. Last accessed: 2025-02-17.
5. Pervez, M.H.; Ecevit, M.I.; Naqvi, N.Z.; Creutzburg, R.; Dag, H. Towards Better Cyber Security Consciousness: The Ease and Danger of OSINT Tools in Exposing Critical Infrastructure Vulnerabilities. In Proceedings of the 2023 8th International Conference on Computer Science and Engineering (UBMK). IEEE, pp. 438–443. https://doi.org/10.1109/UBMK59864.2023.10286573.
6. Weber, S.B.; Stein, S.; Pilgermann, M.; Schrader, T. Attack Detection for Medical Cyber-Physical Systems–A Systematic Literature Review. *11*, 41796–41815. https://doi.org/10.1109/ACCESS.2023.3270225.
7. Berghel, H. Robert David Steele on OSINT. *47*, 76–81. https://doi.org/10.1109/MC.2014.191.
8. Durumeric, Z.; Wustrow, E.; Halderman, J.A. ZMap: Fast Internet-wide Scanning and Its Security Applications. In Proceedings of the 22nd USENIX Security Symposium (USENIX Security 13), Washington, D.C., 2013; pp. 605–620.
9. Aggarwal, M. A Study of CVSS v4.0: A CVE Scoring System. In Proceedings of the 2023 6th International Conference on Contemporary Computing and Informatics (IC3I). IEEE, pp. 1180–1186. Place: Gautam Buddha Nagar, India, https://doi.org/10.1109/IC3I59117.2023.10397701.
10. Heinl, P.; Patapovas, A.; Pilgermann, M. Towards AI-enabled Cyber Threat Assessment in the Health Sector. Version Number: 1, https://doi.org/10.48550/ARXIV.2409.12765.
11. Walkowski, M.; Oko, J.; Sujecki, S. Vulnerability Management Models Using a Common Vulnerability Scoring System. *11*, 8735. https://doi.org/10.3390/app11188735.