In Plain Sight: A Pragmatic Exploration of the Italian Medical Landscape (In)security

Lorenzo Bracciale[1], Pierpaolo Loreti[1], Emanuele Raso[1,*] and Giuseppe Bianchi[1]

[1]*Department of Electronic Engineering, University of Rome "Tor Vergata", Rome, Italy*

Abstract

Protecting the medical sector from ongoing cybersecurity threats poses a highly complex challenge due to its unique combination of highly specialized and domain-specific technologies, coupled with an endemic lack of resources and skill gaps. In assessing the maturity level of Italy's healthcare cybersecurity landscape, we showcase four concrete examples of glaring data leakage and exposed vulnerabilities, illustrating how seemingly trivial issues that could be easily checked or fixed are left unattended. We then offer insights into the reasons behind the occurrence of these basic flaws and suggest alternative strategies that might assist the Italian healthcare sector in addressing cyber threats more effectively, thereby ensuring an adequate level of security to protect health information.

Keywords

Cybersecurity, vulnerability, healthcare

1. Introduction

The ongoing digital transformation is bringing impressive benefits to all societal sectors, but it also comes with new concerns. The medical ecosystem, the focus of this paper, is a prominent example where the convergence of Medical Devices (MDs), Electronic Health Records (EHRs), and interconnected networks has ushered in a new era of efficiency and patient care.

However, this digital transformation is also presenting unprecedented challenges, particularly in securing the vast, heterogeneous, and strictly regulated ecosystem that constitutes the medical landscape. Establishing a robust defense not only necessitates the presence of an IT security team but, at least in principle, should ideally extend to the creation of a comprehensive Security Operations Center (SOC), Cybersecurity Awareness Training, Computer Security Incident Response Team (CSIRT), Health Information Management Team, and Biomedical Equipment Security Team. Furthermore, effective defense requires dedicated Supply Chain and Vendor Management Teams, especially considering the increasing reliance on third-party vendors in the healthcare sector. Compliance with regulations and engagement of Executive Leaders further adds to the multifaceted nature of securing the medical landscape.

But how viable is all of this for a sector like healthcare? One critical aspect contributing to the complexity of medical cybersecurity is the substantial investment required to defend against

ITASEC 2024: The Italian Conference on CyberSecurity

*Corresponding author.

lorenzo.bracciale@uniroma2.it (L. Bracciale); pierpaolo.loreti@uniroma2.it (P. Loreti); emanuele.raso@uniroma2.it (E. Raso); giuseppe.bianchi@uniroma2.it (G. Bianchi)

0000-0002-6673-3157 (L. Bracciale); 0000-0002-2348-5077 (P. Loreti); 0000-0003-1195-6529 (E. Raso); 0000-0001-7277-7423 (G. Bianchi)

potential threats [1]. In the face of such needs, the Italian healthcare ecosystem is currently grappling with a critical shortage of funds and a historical deficit in essential IT skills.

In this context, the core of this paper is dedicated to raise awareness about the insufficient level of cybersecurity hygiene in the Italian medical landscape. Through five pragmatic experimental examples belonging to four different categories (Sections 3-6), and with no pretense of a systematic exploration, we highlight glaring security flaws that should be prioritized before burdening the medical sector with the mandate of deploying more advanced protection solutions. Only by first focusing on the low-hanging fruit — those easily identifiable and fixable issues — can we lay a solid foundation for improvement. This approach is crucial, especially considering that many vulnerabilities highlighted in the subsequent sections stem from elementary and baseline errors, making them relatively straightforward to rectify.

In such a scenario, and keeping in mind the substantial skill gap and resource shortage characterizing the Italian healthcare sector, we hardly believe that further cybersecurity prescriptions will solve the problem. Neither is a "lift and shift" solution, involving the migration to centralized and more easily controllable cloud systems, attainable in the foreseeable future. This approach would face incredible difficulties due to the necessary coexistence with outdated legacy systems and the cyberphysical nature of most medical instruments and devices. A companion goal of this paper is thus to provide our view on practical processes and strategies that can be employed to enhance the cybersecurity health of the medical landscape, in a way that we believe is compatible with the limited resources and skill gaps affecting the medical sector.

The remainder of the paper is structured as follows. After a brief discussion in Section 2 about the Italian healthcare landscape in relation to cybersecurity issues, Section 3 presents two examples of sensitive data currently exposed in plain sight on the Internet. Section 4 demonstrates how potential vulnerabilities can be located within healthcare premises via correlation among public data gathered from various open sources. Hardcoded cryptographic credentials in medical devices are a long-standing problem, exemplified in Section 5 with (yet another) real-world example. More instructive is the incident presented in Section 6. This is a very elementary Stored Cross-Site Scripting (XSS) vulnerability that we spotted months ago over the Italian EHR system and responsibly disclosed to a national CSIRT. This case clearly highlights how a very basic issue may become lengthy and non-trivial to solve when occurring in the highly fragmented Italian medical landscape. While insights from each case are discussed in their respective sections, we introduce an additional Section 7, where we offer our perspective on potential future directions for cost-effective and practical enhancements to cybersecurity in the medical landscape. The paper concludes with some final remarks.

2. Cybersecurity and the healthcare landscape

The ubiquitous incorporation of digital technologies into hospital infrastructures, the capillary distribution of interconnected devices in the emerging paradigm of the Internet of Medical Things (IoMT) [2], and the highly heterogeneous nature of the healtcare systems and MD fabrics comes along with novel cybersecurity threats and vulnerabilities [3, 4, 5].

Securing the medical landscape indeed proves to be a difficult task for several reasons. The very first one resides in the sensitive and critical data handled in the medical context, and its

high impact on people's lives and potential life-threatening implications. Especially the high sensitivity of data often translates into a powerful motivation for cybercriminals to demand ransom for not disclosing the data that has been exfiltrated.

Another unique aspect of the medical landscape resides in the huge attack surface exposed. This is composed not only of cloud servers, gateways [6], and desktop computers, but also includes mobile devices, medical devices [7] and at-home Point-of-Care [8] which are drammatically improving the quality of life of patients with chronic diseases. The heterogeneity of these technologies, ranging from diagnostic equipment to patient monitoring devices, introduces complexities that demand specialized security measures. Additionally, the specialization of protocols and systems, such as Digital Imaging and Communications in Medicine (DICOM), Health Level Seven (HL7) Fast Healthcare Interoperability Resources (FHIR), and Picture Archiving and Communication Systems (PACS), further complicates the security landscape.

Furthermore, the Italian scenario presents unique challenges due to the juxtaposition of modern technologies with outdated infrastructures. According to a 2023 survey conducted by the Digital Health Observatory of the Polytechnic of Milan [9], only 42% of Italian healthcare facilities currently employ Electronic Medical Records (EMRs) comprehensively across all wards, resulting in a 69% of Healthcare Professionnels (HCPs), significantly lower than the European average of 81%. Compounding the issue, Italy's healthcare system operates at a regional level, potentially resulting in up to 20 distinct systems across its regions. This decentralized structure, coupled with endemic resource shortages and ICT skill gaps, may lead to duplicated service implementations and less rigorously reviewed code, thereby increasing vulnerability to exploitation (see Section 6).

At the same time, according to a report by Trend Micro [10] for the first half of 2023, Italy emerged as the most affected country in Europe and the third worldwide by malware. In the first half of 2023, healthcare was the second most targeted sector, with 14.5% of total cases according to a Clusit report [11]. Obviously, in the face of such trends, and with the NIS2 directives deadlines rapidly approaching[1], the Italian healthcare system needs to undertake major investments in cybersecurity [9]. It is less obvious, at least to us, how to best pursue such a goal, and especially whether assigning additional cybersecurity duties to medical infrastructures might be a sound strategy, given the current skill gaps and resource shortage.

3. Publicly Exposed Medical Data: two examples

The combination of practice and skills in using publicly or commercially available search and crawling tools, along with a baseline understanding of medical information technology terms and concepts, allows researchers to discover a significant amount of unintentionally exposed data on the Internet. The following two examples are meant to provide evidence of the degree to which the inadvertent exposure of sensitive information is currently detrimentally impacting the national medical landscape.

PACS servers in plain sight. A plethora of cyberspace-related search services and engines, including Shodan, Censys, Zoomeye, Fofa, NTI [12], and many others, are today available over

[1] https://eur-lex.europa.eu/eli/dir/2022/2555/oj

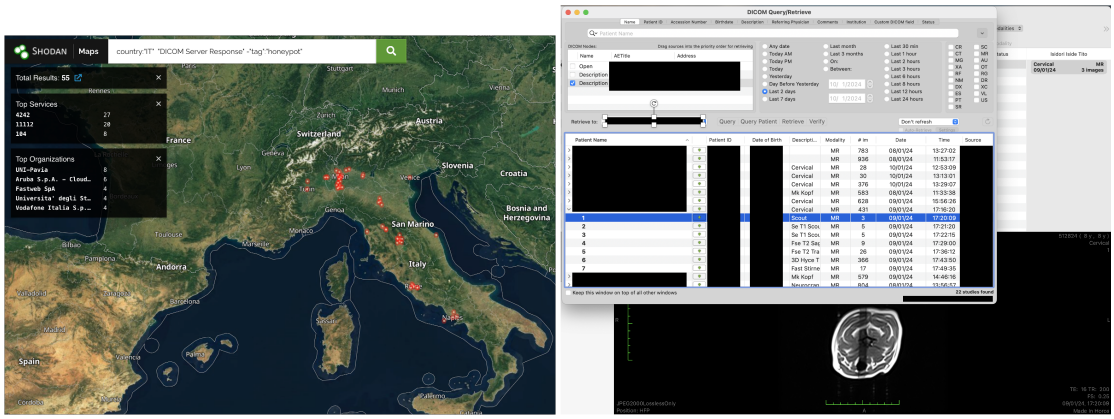

(a) Publicly exposed PACS in Italy

(b) Effect on a potential exposed system

Figure 1: Misconfigured PACSs

the Internet. Among those, Shodan (probably the most known Internet of Things (IoT) search engine [13]) allows to discover and gather information about servers, routers and other Internet-connected devices. Shodan's application to the IoMTs is largely established; for instance, in 2017, a research documented in [14] has revealed over 3,900 vulnerabilities across approximately 1,600 IoT connected medical devices.

Rather than targeting medical devices, we ran Shodan queries to identify servers storing medical data. The majority of healthcare organizations rely on domain-specific servers known as PACS, used to archive medical images and enable HCPs to share these patient records and images with other providers. PACS systems receive images from various imaging modalities such as X-ray, Magnetic Resonance Imaging (MRI), Computed Tomography (CT), ultrasound, and nuclear medicine. The most widely used format for this type of report, called DICOM, is a standard created in 1985 to define the storage and transmission of medical information, with some known security issues [15] that allows the crafting of a malware as a legit DICOM file. DICOM servers can be easily detected as they usually are associated to ports 104 or 11112, and upon a query on these ports, they do respond with the standard string `DICOM Server Response`.

We utilized this method to search for PACS instances exposed on public IP addresses in Italy. The results are depicted in Figure 1a. In the simulated scenario shown in Figure 1b, we illustrate the potential consequences if some of these PACS were left open: any user with a DICOM client could potentially access thousands of medical records without requiring a password. This poses not only a privacy concern but also a security risk, as the permitted DICOM operations include not only `C-FIND` and `C-MOVE` but also `C-STORE`, which enables the *transmission* of DICOM files (potentially containing malware or fake data [15]) to a PACS.

Misconfigured Cloud Buckets. The current surge in healthcare's adoption of cloud technology can be attributed to the growing digitalization of medical applications and the potential improvements in efficiency, cost reduction, and the security and privacy of medical data. In fact, a significant amount of medical data is transmitted through cloud-based medical applications, utilizing services such as AWS S3 buckets, Azure Blob Storage, Digital Ocean Spaces, or Google

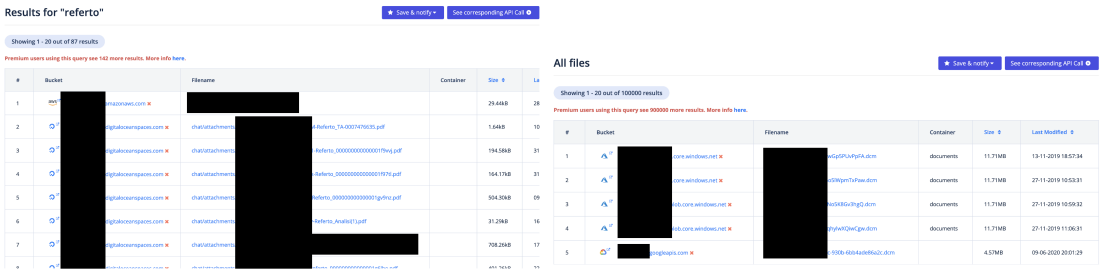

(a) Medical reports exposed (b) 100,000 ".dcm" files exposed

Figure 2: Misconfigured buckets leaks medical data

Cloud Platform.

The data stored in these buckets must, of course, be protected with adequate access control, a common but not always implemented practice. Unfortunately, some buckets expose all data as "public on internet," either due to misconfigurations or the risky practice of protecting files solely by assigning them unguessable names. Confidentiality is compromised in both cases, especially when cloud providers offer a *listing* of files within the buckets. This has led to the emergence of web crawlers like `GrayHat Warfare`[2] or `OpenBuckets`[3].

We employed such search engines to investigate various medical terms, aiming to determine whether these public searches could potentially allow adversaries to intercept inadvertently left unprotected sensitive data. Figure 2a illustrates an example of such a search conducted using the Italian keyword "referto" (i.e., patient medical report). The number of identified records (less than 100 with free `OpenBuckets` access, but increasing to about 200 with premium access) may not be significant by itself, but it serves as clear evidence of the existence of misconfigured buckets storing medical data, potentially totaling millions of data points with different names. Notably, when searching for the ".dcm" file extension (representing DICOM files, i.e., files storing medical data), the results reached 100,000 with free access and increased to about 1 million with premium access (Figure 2b).

Lessons Learned. It is fair to note that the above findings are certainly not new or unexpected. In September 2019, ProPublica revealed that millions of medical images were being exposed online through unsecured PACS. Subsequently, NNT discovered more than 2 petabytes of unprotected medical data on PACS servers, leading to 13 million medical examinations related to approximately 3.5 million U.S. patients being exposed, unprotected, and accessible to anyone on the internet [16].

What surprised us, however, is that, following the initially noticed surge of misconfigured systems, the number of exposed systems actually *increased* over time! Particularly concerning is the fact that this data is in plain sight over the Internet and exploitable even by individuals with no computer skills. Indeed, such analyses are accessible to everyone via trivial-to-use online tools, and do not depend on any specific leak or data breach. These examples illustrate how *reconnaissance*, the initial phase of the kill chain, can be easily conducted in the medical domain,

[2]https://buckets.grayhatwarfare.com/
[3]https://www.openbuckets.io/

highlighting the extensive attack surface of healthcare that inevitably leads to numerous privacy and security concerns - in the words of Mark Dowd, "The attack surface is the vulnerability. Finding a bug there is just a detail".

On the other hand, the natural emerging question is: *if such reconnaissance is so easy, why is it not systematically conducted - for prevention purposes - by our national authorities?* In fact, while detailed reports in this area are found in other countries, we observed a dearth of specific reports tailored to the Italian market.

4. Localizing Vulnerable Medical Devices

While the elementary queries shown in the previous section were based on public engines, the more "creative" use of open data highlighted in what follows permits to gather further evidence about the presence of potentially vulnerable MDs *inside* specific Italian hospitals.

For transparency purposes, in Italy (but also in many other countries, see [7]), all purchases made by the public administration must be traceable. The corresponding list, detailing which institution purchased which asset, is made public and accessible to anyone by the National Anti-Corruption Authority (ANAC). Meanwhile, the products, systems, and medical devices listed in these purchases may, sooner or later, be affected by vulnerabilities. When a new flaw in the software of a medical device is disclosed, similar to any other kind of software, it is reported in the Common Vulnerabilities and Exposures (CVE) worldwide database, public as well. It follows that, by cross-referencing the purchase logs provided by ANAC with the vulnerability databases, *anyone can determine the specific hospital or healthcare facility where a potentially vulnerable medical device is located.*

By applying a methodology proposed in our previous work [7] to the Italian scenario, we examined purchase orders issued by the Italian public administration up to mid-2022. Using data mining techniques, we identified purchases of MDs known to have vulnerabilities, as indicated by cybersecurity alerts issued by the US's Cybersecurity and Infrastructure Security Agency (CISA). We define a *match* as a situation where we can reasonably attribute a purchase to a potentially vulnerable device, as shown in Figure 3. This does not only refer to the acquisition of a single device but may also include the procurement of spare parts or consumables indirectly indicating the presence of the device within a specific healthcare facility.

Our analysis establishes a detailed timeline of when healthcare facilities acquired potentially compromised devices. It is important to note that the presence of these devices in purchase records does not necessarily imply that they remain vulnerable at present. Many manufacturers proactively recall MDs that pose health risks. However, the recall process may encounter challenges related to the effectiveness of communication among manufacturers, healthcare providers, and patients [17].

The specific results of such an analysis, tailored to the Italian scenario and reported in the Appendix, underscore that the diffusion status of potentially vulnerable devices within healthcare facilities appears quite critical, either in terms of exposure time (the time between the purchase of the device and the publication of the vulnerability, more than 3 years on average) as well as in terms of severity of the vulnerabilities, measured via their CVSS score.

Lessons Learned. While, at first glance, the described approach might be viewed as susceptible

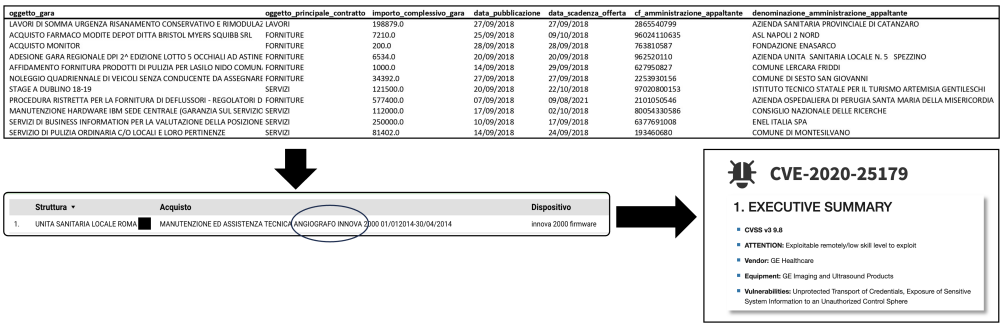

Figure 3: Example of a *match*: a purchased medical device is linked to a CVE

to adversarial behavior, allowing the identification of vulnerable devices within healthcare institutions, we believe it serves as a compelling example of how systematic correlation among (open) data, conducted for preventive purposes, can effectively manage risks and enhance awareness. We would actually advocate the extension of the described approach to incorporate additional non-public data, such as that collected by hospitals' inventory systems, thereby serving as valuable resources for authorities aiming to improve the cost-effectiveness of security audits.

5. Hardcoded Credentials in Medical Mobile Applications

According to the World Health Organization (WHO), there are more than 2 million different types of MDs [18]. Many of these devices, such as insulin pumps, not only communicate with mobile applications to display parameters (e.g., blood glucose levels) but also perform actions (e.g., control insulin delivery and administer boluses). Some of these systems, however, have been found to be severely insecure for various reasons. Many custom-made protocols either lack proper authentication or authorization mechanisms, or fail to implement them effectively [19, 20], exposing vulnerabilities that could result in potentially fatal consequences [21, 22]. In other cases, although authentication and data protection measures are theoretically in place, the relevant credentials and/or secrets are hardcoded in the software. Unfortunately, this is a fairly common poor practice in medical apps, as highlighted by Alissa Knight in her well-known white paper [6].

With no pretense of a systematic analysis, we here limit to show yet another contemporary example of hardcoded credentials (kept anonymous for obvious reasons, and being just one among many other similar cases). This example refers to a MD controlled via a mobile app. In this case, code inspection is especially easy as we can directly analyze the app code, without the need to extract the device's firmware.

As a first step in our analysis, we employed an automatic static Java code analyzer (MobSF), uncovering no specific warnings. However, from a following manual analysis we found some interesting elements not noticed by the automatic analyzer.

In a medical mobile app, the analyzer overlooked a shared secret key, as shown in Figure 4a. This likely occurred because it attempted to match strings, whereas the key was represented as an array of bytes.

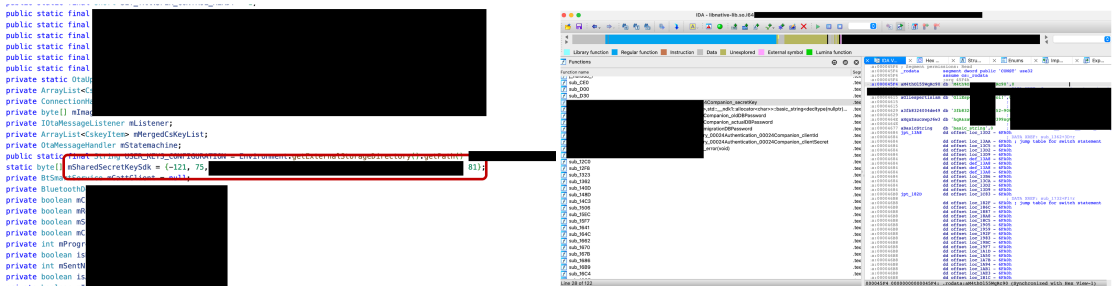

(a) Hardcoded Shared Secret Key (b) Hardcoded credentials in native libraries

Figure 4: Example of hardcoded credentials in mobile applications connected to medical devices

In another medical app we found an insecurity which could be releated with the recognized trend in current mobile app development to move away from Java classes and Android DEX files. To improve performance and access platform-specific features, developers are increasingly opting for the use of "native" code or libraries, which directly interact with the device's hardware and operating system. Native code, typically written in languages like C or C++, is compiled and loaded onto the mobile device as binary.

We cannot determine whether the dependence on native code might instill developers with a falsely perceived sense of enhanced security. However, it is evident, as illustrated in Figures 4b, that in our specific case, a native component was utilized to store multiple hardcoded credentials. Clearly, this is easily discoverable with minimal familiarity with ELF files and a good disassembler, allowing the exposure of credentials[4] in plain sight as demonstrated in Figure 4b.

Lessons Learned. Since reverse engineering is today affordable to any practicioner, there is little more to add to this example beyond quoting the renowned computer security expert Thai Duong: "Assume that your opponents know everything: if the source code or design blueprint were leaked today, it should not change the security exposure of your product."

6. Electronic Health Records: A Tale of a Vulnerability

EHRs are centralized, patient-focused repositories accessible immediately and securely to permitted individuals, predominantly physicians. One of the key difference with Electronic Medical Records (EMRs), is that the core purpose of EHRs is to facilitate the *exchange of data*, governed by interoperability standards such as HL7, IHE or LOINC. Within the context of Europe's digital health strategy, EHRs stand as a fundamental component. The standardized format for EHR exchange is designed to enable citizens to effortlessly obtain and disseminate their medical information among healthcare providers, including when seeking specialized care or confronting medical emergencies throughout the EU. The italian EHR is called "Fascicolo Sanitario Elettronico" (FSE), and was instututed in 2012 by Decree-Law No. 179, with a collaboration

[4]We did not further explore the impact of such credential leakage, i.e. whether these are related to significant or harmful tasks or are just left over from some obsolete function, i.e., just as a legacy of a bad practice.

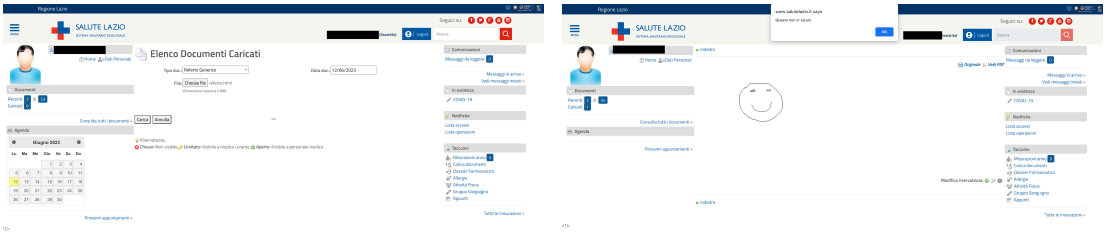

(a) Uploading malicious data in Patient-Originated Data

(b) Execution example of the stored Cross-Site Scripting

Figure 5: Insecurity on EHR patient-generated data with stored Cross-Site Scripting

between the Agency for Digital Italy (AGID), Ministry of Health and Ministry of Economy, together with the Data Protection Authority and the Regions.

Technically, the italian EHR has a complex structure, ruled by many standards on content (e.g., HL7 CDAR2), interoperability profiles (IHE ITI), and functional requirements (HL7 EHR-S). In practice, the resulting structure today is a big federated data storage. The access to such data is demanded to specific applications such as commercial medical software which are connected through APIs, or to public website provided for free by the public administration to give a free access to citizens and, at times, HCPs.

The implementation of the EHR is demanded to the 20 Italian Regions, giving only the functional requirements defined by the HL7 EHR-S FM specification, so that each region is called to basically implement the same thing in a potential different way. Today, the National Recovery and Resilience Plan (PNRR, NextGenerationEU), in its Mission 6, put a specific goal of reaching a 85% of utilization by general practitioners before 2025 causing a certain rush necessary to reach the deadline. In the meanwhile, the new version "2.0" of the Italian EHR is moving its first step with the goal of speeding up and improving the digitalization, adding new features and elements such as validation Gateways and a Central Data Repository.

Code duplication, sensible data, rush, new features: this is a fertile soil for the growth of cybersecurity problems. The Italian Data Protection Authority brought one of the first case to the public's attention with case number 9883731 on March 23, 2023. The EHR provider for the Autonomous Province of Bolzano was subjected to a fine of €15,000. The penalty was imposed due to the provider's failure to establish adequate authorization controls within their system. This significant oversight allowed individuals to access their own EHR and then alter the `patient_id` field in the URL. By inputting another person's tax code, they could unlawfully view the health data of other individuals.

We spot another insecurity in a different implementation of the EHR. Specifically, we analyse a component of EHR called *Patient-Originated Data* ("Taccuino"). This component is designed for patients to input their own health information, which may stem from personal recollections or notes transcribed from physical documents. Unlike other parts of the EHR, this data is not introduced by a physician or an HCP but is directly entered by the patient, which means by any Italian citizen. This particular feature presents a security risk by potentially allowing the upload of malicious data that could be accessed by physicians. More specifically, it opens the door for

stored XSS attacks, especially when combined with the EHR public web interface. Through this method, a patient could upload a file containing malicious "HTML+JavaScript" code, for instance disguised as a standard medical report, as depicted in Figure 5. Any doctor opening such data in their browser will execute the JavaScript code, leading to potential unauthorized actions on the EHR portal, such as data theft, unwanted actions on behalf of the phisician, or credential theft through hardly recognizable XSS-based spear phishing schemes. The issue's severity is amplified by the shared nature of EHR data, since EHR is not a single software or a single website. It is thus needed a thorough examination of *all* applications with EHR access to ensure there are safeguards against *both* the upload and download of malicious data.

The vulnerability was notified to CSIRT for a Coordinated Vulnerability Disclosure and apparently the problem has been resolved after *7 months*.

Lessons Learned. The reported incident provided us with two crucial insights. Firstly, we discovered that the resolution times for bugs in the Italian medical landscape can be exceptionally lengthy. Although this phenomenon has been extensively documented in the literature, particularly in many U.S. cases, experiencing and quantifying it firsthand in our national context was enlightening. The prolonged timeframe results from both the inherent difficulty of implementing changes to such sensitive systems in a production environment and the intricate nature of the healthcare ecosystem, involving multiple stakeholders.

Secondly, each Italian region autonomously applied a patch to address the identified vulnerability. Unfortunately, no information was provided regarding how the effectiveness of the patch was verified. Regrettably, due to the lack of authorization for further experiments, we cannot assess whether all regions have correctly implemented an appropriate patch. It's important to note that type checking in file uploads is a nuanced task, considering the diverse techniques employed by penetration testers to circumvent common defense mechanisms.

7. Discussion and suggestions

In this section we take the freedom to propose three suggestions for improving the security posture of our medical systems.

1. Perform centralized continuous active monitoring with reconnaissance tools. Reconnaissance is usually the initial step in the Cyber Kill Chain, and numerous commercial products (e.g., Shodan, Censys, Zoomeye, Fofa, NTI, etc.) exist, which can be used to gather information about potential data leakage, misconfigured cloud applications, or exposed PACS. As demonstrated in Section 3, these same tools can also be employed for discovering security or privacy issues. In fact, such systems do even more. They can also recognize potential vulnerabilities in medical services, correlating the "banners" with well known CVEs [14, 23, 12]. Instead of relying on each healthcare institution to deploy its individual monitoring tools, we strongly recommend their centralized adoption by one or more national associations or agencies. Not only would this represent an extremely cost-effective and swift solution for the continuous monitoring of our most critical medical systems (e.g., PACS, HL7 Gateways, and Radiotherapy Systems [24]), but centralized management would ensure a consistent and secure configuration of the tools and the relevant applied queries/policies. This model is already implemented in some other countries; for example, Health Information Sharing and Analysis

Center (Health-ISAC) already provides monitoring solutions for free to all its members, along with producing aggregated reports on the state of cybersecurity in healthcare [25].

2. Changing approach: from prescriptive to supportive. Many of the vulnerabilities identified, such as those in Section 3, have straightforward solutions. An alert from an authority to developers/maintainers would be sufficient to implement a quick fix, such as utilizing pre-signed URLs for open buckets or setting a password to enhance security in the case of exposed PACS. These are instances of specific, concrete cases that warrant direct notification. However, the observed Italian trend is quite the opposite, and we are afraid that there is a risk of drifting towards an overly prescriptive approach. The systematic transferring of *any* cybersecurity problem directly onto healthcare institutions can be not only overwhelming and very costly, but also ineffective in front of a gap in specialized cybersecurity expertise. In addition, most of the current awareness campaigns revolve around generic issues (e.g. phishing). This poses a challenge in information management, as an excess or overly broad array of information can hinder the capacity to focus on pertinent and significant elements. Conversely, stakeholders in the medical domain would benefit of advice tailored to their specific cases, and of comprehensive yet focused information, catering to both managerial and non-technical personnel. This can be achieved partially at a low cost and through automated systems, as demonstrated in Sections 3 and 4. Finally, to assist the public sector, there is a need for cross-departmental technical teams of security experts who can collaborate and test systems alongside developers following modern business practices. This would ensure that cases like the one reported in Section 6 receive an effective fix.

3. Encouraging community involvement. Especially in these times of skill shortage, the active involvement of the cybersecurity community could significantly boost the identification of potential vulnerabilities and threats, and even help ensuring a more comprehensive and effective bug fixing. Indeed, active involvement of volunteers has been already started in several international scenarios. For instance, health portals such as DoctoLib (300,000 HCPs) incentivize with monetary rewards (up to €25,000) those who report a vulnerability, through Bug Bounties. They also offer the full list of their APIs on their site to facilitate the work of testers, and clearly write the scope and the rules of engagement. Even public organizations such as the WHO offer incentives for reporting bugs (e.g., publication in a hall of fame), giving details on qualifying vulnerabilities and on reporting rules. In the healthcare sector, which is predominantly public in Italy, there is a need to follow these practices, incentivizing Coordinated Vulnerability Disclosure and bug reporting for example with economic (like DoctoLib) or social credits (like WHO). Public scrutiny and extensive security testing of interfaces and products are arguably the best tools we have to strengthen the defenses of our medical systems. Finally, open source development and community efforts should be promoted, and we commend the Ministry of Health for actively developing part of the "FSE 2.0" on GitHub.

8. Conclusion

Through this paper, we have presented several experimental examples aimed at highlighting the critical cybersecurity status of the Italian healthcare landscape. A concerning aspect is that most of the security flaws we emphasize can be identified using readily available Internet tools,

requiring no specific expertise. While these issues are fixable, we argue that imposing excessive cybersecurity provisions on healthcare institutions can be overwhelming, costly, and ineffective, especially given the shortage of specialized expertise. Consequently, we advocate for strategic improvements, such as the centralized adoption of continuous active monitoring tools, and an increased engagement of the cybersecurity community to enhance vulnerability identification and facilitate more comprehensive and effective bug-fixing initiatives.

Acknowledgments

This work has been partially funded by the Rome Technopole Project (PNRR - NextGenerationEU).

References

[1] A. J. Cartwright, The elephant in the room: cybersecurity in healthcare, Journal of Clinical Monitoring and Computing (2023) 1–10.

[2] S. Razdan, S. Sharma, Internet of medical things (iomt): Overview, emerging technologies, and case studies, IETE technical review 39 (2022) 775–788.

[3] M. Papaioannou, M. Karageorgou, G. Mantas, V. Sucasas, I. Essop, J. Rodriguez, D. Lymberopoulos, A survey on security threats and countermeasures in internet of medical things (iomt), Transactions on Emerging Telecommunications Technologies 33 (2022) e4049.

[4] M. K. Hasan, T. M. Ghazal, R. A. Saeed, B. Pandey, H. Gohel, A. Eshmawi, S. Abdel-Khalek, H. M. Alkhassawneh, A review on security threats, vulnerabilities, and counter measures of 5g enabled internet-of-medical-things, IET Communications 16 (2022) 421–432.

[5] G. Hatzivasilis, O. Soultatos, S. Ioannidis, C. Verikoukis, G. Demetriou, C. Tsatsoulis, Review of security and privacy for the internet of medical things (iomt), in: 2019 15th international conference on distributed computing in sensor systems (DCOSS), IEEE, 2019, pp. 457–464.

[6] A. Knight, Playing with fhir: hacking and securing fhir apis, 2020.

[7] L. Bracciale, P. Loreti, G. Bianchi, Cybersecurity vulnerability analysis of medical devices purchased by national health services, Scientific Reports 13 (2023) 19509.

[8] G. M. Bianco, E. Raso, L. Fiore, V. Mazzaracchio, L. Bracciale, F. Arduini, P. Loreti, G. Marrocco, C. Occhiuzzi, Uhf rfid and nfc point-of-care–architecture, security, and implementation, IEEE Journal of Radio Frequency Identification (2023).

[9] TechFlix360, Il sistema sanitario italiano è a un momento di svolta: l'analisi dell'osservatorio sanità digitale, 2024.

[10] T. Micro, Stepping ahead of risk, 2023. https://d110erj175o600.cloudfront.net/wp-content/uploads/2023/09/12141231/Infographics-TREND-MICRO-2023-MIDYEAR-THREAT-REPORT.pdf.

[11] Clusit, Rapporto clusit 2023 sulla sicurezza ict in italia, 2023. https://clusit.it/wp-content/uploads/download/Rapporto_Clusit_aggiornamento_10-2023_web.pdf.

[12] B. Zhao, S. Ji, W.-H. Lee, C. Lin, H. Weng, J. Wu, P. Zhou, L. Fang, R. Beyah, A large-scale empirical study on the vulnerability of deployed iot devices, IEEE Transactions on Dependable and Secure Computing 19 (2020) 1826–1840.
[13] A. Albataineh, I. Alsmadi, Iot and the risk of internet exposure: Risk assessment using shodan queries, in: 2019 IEEE 20th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM), 2019.
[14] E. McMahon, R. Williams, M. El, S. Samtani, M. Patton, H. Chen, Assessing medical device vulnerabilities on the internet of things, in: 2017 IEEE international conference on intelligence and security informatics (ISI), IEEE, 2017, pp. 176–178.
[15] M. Picado Ortiz, Hipaa-protected malware? misusing dicom flaw to embed malware in ct/mri imagery, 2019.
[16] D. Schrader, Cybersecurity threats in us healthcare systems exposed, 2020.
[17] R. Zipp, Anatomy of a medical device recall: How defective products can slip through an outdated system, Available online., 2021. https://www.medtechdive.com/news/medical-device-recall-process-fda-philips-medtronic/608205/ (visited: 2023-09-14).
[18] W. H. Organization, World health organization - medical devices, Available online., 2023. https://www.who.int/health-topics/medical-devices (visited: 2023-05-20).
[19] NVD, CVE-2019-10964., Available from MITRE, CVE-ID CVE-2019-10964., 2019. http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-10964 (visited: 2023-05-20).
[20] U. P. A. Networks, Know Your Infusion Pump Vulnerabilities and Secure Your Healthcare Organization, Technical Report, Unit42, 2022.
[21] NVD, CVE-2021-42744., Available from MITRE, CVE-ID CVE-2021-42744., 2021. http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2021-42744 (visited: 2023-05-20).
[22] W. Saltzstein, Bluetooth wireless technology cybersecurity and diabetes technology devices, Journal of diabetes science and technology 14 (2020) 1111–1115.
[23] C. Tziampazis, Exposure Assessment on Medical Devices in the Netherlands, B.S. thesis, University of Twente, 2019.
[24] F. R. Labs, The enterprise of things security report the state of iot security, 2020.
[25] Health-ISAC, State of cybersecurity for medical devices and healthcare systems, Available online., 2023. https://h-isac.org/2023-state-of-cybersecurity-for-medical-devices-and-healthcare-systems/ (visited: 2023-09-14).

A. Appendix A: Statistics of Vulnerable Medical Devices

Figure 6a illustrates the number of matches concerning the year of purchase. The result is a detailed presence map of potentially vulnerable devices purchased by italian healthcare facilities. Figure 6b shows the time between the purchase and the publication of a vulnerability of a given device. This *exposure window* represents the potential time for which we have vulnerable devices in our medical facilities, which is 3.5 years on average. It is interesting to show also the severity of the interested vulnerability expresses with their CVSS Score in Figure 7, depicting a landscape where easy-to-exploit vulnerabilities result in a severe impact on the confidentiality, integrity, and availability of medical devices. Finally, we show in Figure 7b and 7c how such

vulnerabilities affect all types of medical devices across the board and involve all MDR risk classes.

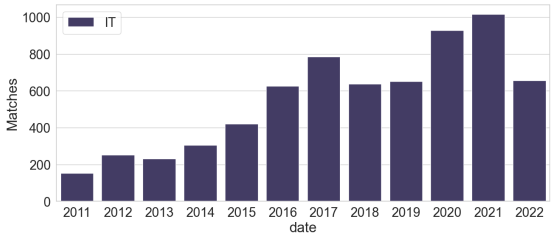

(a) According to the year

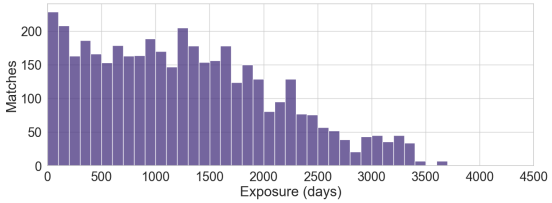

(b) According to the exposure time

Figure 6: Medical devices statistics

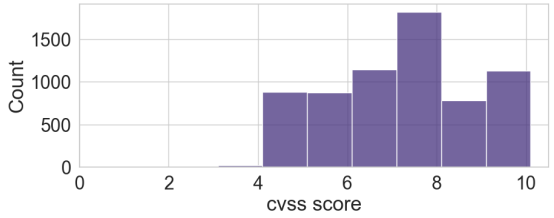

(a) According to the number of matches

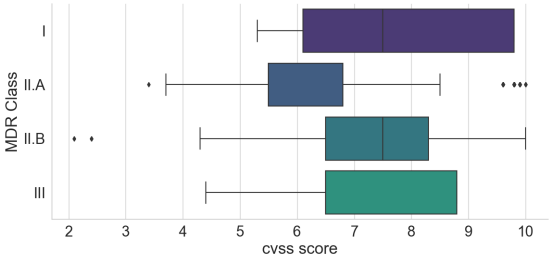

(b) According to the device class

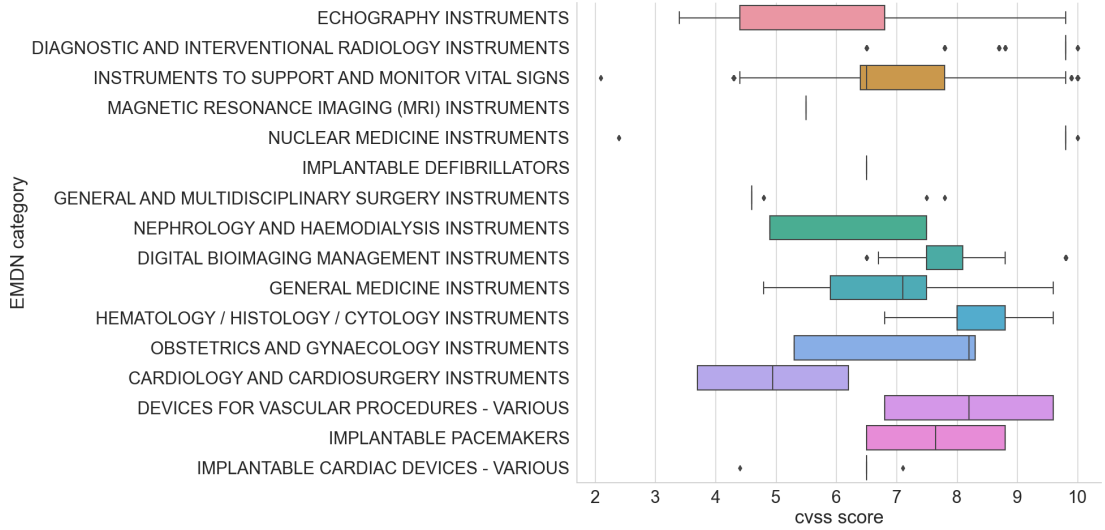

(c) According to the device EMDN category

Figure 7: Analysis of CVSS score