2021 13th International Conference on Cyber Conflict
Going Viral
T. Jančárková, L. Lindström, G. Visky, P. Zotz (Eds.)
2021 © NATO CCDCOE Publications, Tallinn

Epidemic? The Attack Surface of German Hospitals during the COVID-19 Pandemic

Johannes Klick
Alpha Strike Labs
Berlin, Germany
j.klick@alphastrike.io

Robert Koch
Universität der Bundeswehr
Neubiberg, Germany
robert.koch@unibw.de

Thomas Brandstetter
Limes Security/FH St. Pölten
Hagenberg, Austria
tbr@limessecurity.com

Abstract: In our paper, we analyze the attack surface of German hospitals and healthcare providers in 2020 during the COVID-19 pandemic. A primary analysis found that 32 percent of the analyzed services were vulnerable to various degrees and that 36 percent of all hospitals showed numerous vulnerabilities. Further resulting vulnerability statistics were mapped against the size of organization and hospital bed count. The analysis looked at the publicly visible attack surface utilizing a Distributed Cyber Recon System, through distributed Internet scanning, Big Data methods, and scan data of almost 1.5 TB from more than 89 different global Internet scans. From the 1,555 identified German hospitals and clinical entities, analysis of the external attack surface was conducted by looking at more than 13,000 service banners for version identification and subsequent CVE-based vulnerability identification.

Keywords: *Distributed Cyber Recon System, German hospitals, cybersecurity, vulnerabilities, attack surface, critical infrastructure*

1. INTRODUCTION

In October 2020, US-CERT issued a warning regarding the increasing ransomware activity in the healthcare sector [1]. It was common knowledge [23] that healthcare organizations were promising targets for ransomware gangs. Surprisingly, at the very beginning of the COVID-19 pandemic, several ransomware gangs actually pledged not to hit hospitals because of the ongoing scourge. The Maze and DoppelPaymer groups, for instance, said they would not target healthcare facilities and, if they accidentally hit them, would provide the decryption keys at no charge. As another example, the Netwalker operators stated they would not intentionally target hospitals; however, if accidentally hit, the hospital would still have to pay the ransom. Unfortunately, other attacker groups did not have such scruples. Ransomware incidents against hospitals skyrocketed in October 2020, most notably with the use of Ryuk ransomware against 250 U.S.-based hospitals and clinics [20]. The criticalness of the ransomware attack wave against the U.S. was demonstrated by the very rare tri-agency ransomware alert issued by the Federal Bureau of Investigation (FBI), the U.S. Department of Health and Human Services (HHS), and the Cybersecurity and Infrastructure Security Agency (CISA), and hosted by the aforementioned US-CERT.

Naturally, in an increasingly digitized and interconnected world, those issues are not limited to the United States. In Germany in 2020, an intense discussion was prompted by an incident involving the death of a patient who had to be taken to a distant hospital because the closest hospital was signed out of emergency treatment due to an ongoing ransomware attack (see, e.g., Ralston [22]). Even though the longer ride to the more distant hospital was later found not to have been a factor in the patient's death, this specific example underscores the increasing threats posed by cyber attacks, particularly in the healthcare sector.

It must be noted, however, that cybersecurity threats in the healthcare and medical sector are anything but new. On the one hand, healthcare and medical production has always been an innovative field, in which new procedures and technologies are used. On the other hand, there are known challenges – specifically, the long life cycles, or rather the long service life of products, in this area, as well as the need for time-consuming re-certifications, such as when changing or patching the software. The need for comprehensive quality control and certification, especially in the medical field, is illustrated by the example of Therac-25 and the fatal incidents involving the faulty irradiation of patients in the 1980s [14]. Although the healthcare equipment of several vendors has a higher security level nowadays, many healthcare components and systems still have numerous security issues, some of which are even critical, according to the Common Vulnerability Scoring System (CVSS), which "provides a way to capture the principal characteristics of a vulnerability and produce a numerical score reflecting its severity" [5]. To make matters worse, the attack surface (vulnerabilities and starting points for an attack) stemming from complex healthcare networks and equipment has become increasingly challenging to defend [13].

Given this increase in cybercrime, the question arises as to what the cybersecurity situation is in the healthcare sector, which weaknesses and vulnerabilities can be identified in the healthcare system infrastructure, and what recommendations for action can be derived from these. In the very topical context of the COVID-19 pandemic, we therefore chose to examine the cyber attack surface in terms of visible vulnerabilities of hospitals and clinical providers in our home country, Germany.

The rest of the paper is structured as follows: the chapter after the introduction gives context – it describes related work and background information on ransomware attacks in the healthcare sector, as well as the overall development of this problem. Chapter 3 describes the technical infrastructure that made our analysis possible: we describe our Distributed Cyber Recon System and how we used and extended it through our analysis. In Chapter 4 our methodology for attack surface detection of hospitals and clinical providers is presented, showing how we approached this from a healthcare entity identification point of view, as well as an attack surface correlation point of view. Chapter 5 contains the data section, where we describe the results of our findings in detail both verbally and graphically. Chapter 6 summarizes the results of our analysis, as well as its shortcomings and our ideas for future work.

2. ON THE HISTORY OF RANSOMWARE

As a result of innovation and (at that time generally) low security standards, the very first piece of ransomware, surprisingly, emerged from the medical sector. In 1989 the malware "PC Cyborg," also commonly known as "AIDS Trojan" [3], was distributed to an estimated 20,000 people, including the participants of a WHO conference on AIDS. Hidden under the guise of evaluation software, the first encryption Trojan was released; it was attributed to the American biologist Dr. Joseph Popp. Interestingly enough, the damaging effects of the Trojan were stated in the user agreements and had to be accepted by the user upfront.

Though this type of malware appeared early in computer history, it took ransomware a long time to achieve "success." Ransomware variants such as "Fake Antivirus" (2001), GPCoder (2005), CRYZIP (2006), and QiaoZhaz (2007) appeared from 2001 onwards, but the attacks were still limited, mainly due to technical reasons and/or logistical obstacles to transferring the money. Some creative approaches like WinLock used SMS and phone calls to premium numbers, for example, to monetize attacks, but a noteworthy crime breakthrough came with CryptoLocker in 2013, introducing payments via Bitcoin. While CryptoLocker was taken down in June 2014, it was the blueprint for numerous copycats, as it showed that it was possible to earn millions within a few weeks. Thus the right combination of public-key cryptography, the digital currency Bitcoin, anonymization possibilities by using the Tor network, and providing a reliable decryption opened up a new criminal business model which today accounts for a large share of the total damage in the billions per year. Fiscutean [6] gives an overview of the history and other details of ransomware.

On the basis of various technical developments and improvements, it was thus possible for criminals to implement an effective digital extortion model by means of simple cryptography, anonymous communication, and straightforward, quasi-anonymous payment options. True, there have been cases in which the data of the attacked system has been destroyed and actual recovery was never intended (for example, because no required key material was kept). However, these were exceptions and stemmed either from errors in technical implementation or simply from the attacker having other intentions than demanding ransom. The success was based on the fact that victims who choose to pay have a good chance of recovering their data; the attackers are thus motivated to enable correct decryption in order to keep their business model alive and thriving.

In the earlier days, the attackers struck out at random; the victims were often individuals and typically could pay only small ransoms. However, over time, the attackers grew more professional. They began targeting large organizations, and their attacks and ransom demands became bolder [18]. Companies active in the grey area, which sell vulnerabilities including 0days (vulnerabilities that are still unknown to the manufacturer of the product), extend this threat. An example of this is the “MedPack” of the company GLEG Ltd., which contains 0days specifically for medical software [15].

The amount of the ransom is, for obvious reasons, based on a corresponding analysis of the target. Blackmailers have also increased the pressure on the victims by threatening to publish data stolen from the company. On several occasions, they have followed through [11].

In theory, this trend can only be interrupted if no more payments are made for a long period of time. The technical prerequisite and basis for this are regular offline backups, as well as regular tests of the disaster recovery procedures, with dedicated checks on ransomware recovery.

Unfortunately, theory and practice are often worlds apart. As Goodwin and Smith [25] found, only half of all apps are fully covered by a disaster recovery strategy. In some cases, backups are not current and up-to-date or just not available due to misconfiguration, perhaps because they are not kept offline and were also encrypted. In other cases, critical aspects of the recovery process fail because they were never validated in the current environment.

Driven by the increase in ransomware attacks, companies are considering either investing in cyber insurance in order to cushion potential financial damage, or, if the situation arises, simply paying the ransom. The statistics are telling: over 40 percent of cyber insurance claims now involve ransomware [4]. Accordingly, some countries are considering banning the payment of ransoms in order to undercut the criminals' business model. The U.S. Department of the Treasury has already pointed out that ransom payments to groups or organizations on the sanctions list are punishable if they are not approved [24] by the Office of Foreign Assets Control (OFAC). "Cyber-related Sanctions" is a special section on the U.S. Department of the Treasury's website.

The difficulty of implementing this requirement is, however, evident from past cases such as the attacks on police departments in Swansea, Massachusetts [17], and in Dickson, Tennessee. These departments, infected by the ransomware CryptoWall 2.0, paid a ransom to recover their data. With this background, it is worthwhile to explore the attack surface and security posture in the healthcare sector.

Given the increase in ransomware campaigns, the outstanding importance of a functioning healthcare system – especially in the prevailing COVID-19 pandemic – and possible influencing factors through the short-term provision and integration of remote access and teleworking possibilities, we have conducted an in-depth investigation of the cyber attack surface of German hospitals based on the "Deutsche Interdisziplinäre Vereinigung für Intensiv- und Notfallmedizin" (DIVI) intensive register [9]. The DIVI register is a national registry of intensive care capacities detailing available and overall intensive-care capacities in Germany. It was created in response to the COVID-19 pandemic.

3. INTRODUCTION TO THE DISTRIBUTED CYBER RECON SYSTEM (DCS)

The previous chapters have illustrated that it is both possible and feasible to attack hospitals and medical devices. However, the question arises: just how large is the potential cyber attack surface of critical infrastructures like hospitals?

Reconnaissance and, in particular, representation of an organization in cyberspace is a major challenge. For this reason, there was a need for a novel search engine that could search the Internet (2.8 billion routed IPv4 addresses) in a few hours for a network service or services/servers with a specific vulnerability in a matter of hours, and which would also allow mapping to a specific target organization.

In our Distributed Cyber Recon System (DCS) developed specifically for various reconnaissance and analysis tasks, we can answer questions like: What is the security posture of a specific organization? What is the attack surface of an entire group of organizations? Which systems belong to which organization in the first place? Since plain Internet scan data is not sufficient, the scan data is augmented with additional information such as WHOIS data, IP prefix, Autonomous System (AS) information, certificate information, and geo-information about the IP of the system. The combination of this information in a Big Data approach enables a quite accurate representation of cyberspace.

For example, the distribution of selected system versions of a particular network service in an organization, as well as all detected Industrial Control Systems (ICS), can be displayed on a map, and systems can be organized by, e.g., specific country. IP prefix and IP ownership information can also be selected and aggregated using dynamic charts. This allows a recon analyst to get a quick overview of their own cyber infrastructures, as well as those of foreign states, organizations, and companies.

In our case, this DCS was used to analyze the publicly visible system attack surface of hospitals located in Germany. In the following passages, the methodology of our data collection, and that of the DCS, is explained in more detail.

The DCS searches the Internet globally from 1,024 different IPv4 addresses. First, TCP SYN scans are performed for 2.8 billion IPv4 addresses. For each response to this scan, the corresponding application protocol, such as HTTPS or Telnet, is scanned. Then, for each IP address, the result is enriched with owner information (Autonomous System (AS) Information), holder information (WHOIS database), Geo-IP data, BGP information, and other data sources from the Internet. For the analysis of a target entity, all data fields in the scan data are searched for the name or domain of the entity. The identified IP addresses and reachable network services are then used for further analysis. Consequently, the DCS always scans the entire Internet and only identifies the associated network areas and network services of a target entity in a post scan phase. This means the DCS database essentially holds the same type of data for any conceivable target set. In the next stage, this target information is made available in a user interface called Inspector for further advanced analyses such as vulnerability matching based on the service banner, subdomain identification, and screenshot generation.

The DCS primarily consists of the search nodes, a backend, and a frontend. The relationships between the individual components are shown in Figure 1. The frontend is used by an analyst for operation setup and data analysis. The IPv4 network ranges, protocols, ports, and scan algorithms to be scanned are defined in the frontend.

FIGURE 1: DISTRIBUTED CYBER RECON SYSTEM ARCHITECTURE

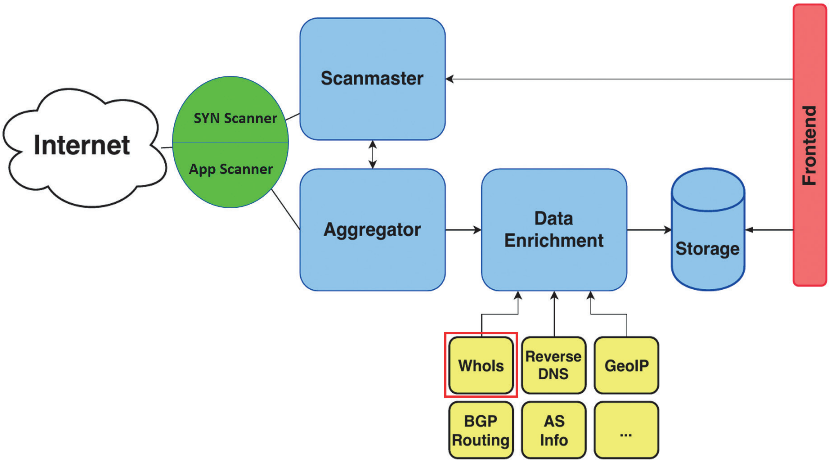

The IPv4 range to be scanned is then pseudorandomized by the scan master in accordance with the selected scan algorithm, divided into several work units and distributed to the various scan nodes. In addition, this measure helps to stay below the triggers of intrusion prevention systems, because the scan traffic is distributed to as many different target networks as possible at the same time. The scan nodes are distributed worldwide for quality and correlation reasons and have different scan bandwidths.

Our DCS enables us to scan the same target areas simultaneously from different strategically interesting locations (e.g., different countries) as site groups and to compare the results. With the help of well-chosen scan locations, potential national Geo-IP blocks can be detected and subsequently bypassed. Experience has shown that result quality can significantly improve with a globally distributed group of scan nodes, as not all destinations are visible from all parts of the Internet due to various national or regional filtering approaches. Furthermore, if a scan node fails, the scan master will automatically detect this. Subsequently, the scan master will assign the work unit of the failed node to a new search node. This ensures that all required IP addresses are always scanned, guaranteeing that the system produces consistent data.

The search nodes consist of two primary components. First of all, the SYN scanner is active, which only sends TCP SYN or UDP packets. During the sending process, the last used destination addresses are stored in a ring buffer. At the same time, the scanner is waiting for incoming TCP SYN ACK packets or UDP responses whose senders correspond to the destination addresses of the ring buffer. This prevents the search engine from being used as a DDoS amplifier. The ring buffer ensures that the search engine only responds to TCP SYN-ACK packets that it has sent out itself. Without the ring buffer, an attacker could send spoofed TCP SYN-ACK packets via IP spoofing, and the search engine would send additional application protocol level scan traffic to the spoofed IP addresses, thus using the search engine as a DDoS amplifier. Furthermore, the *search nodes use a total of more than 1,024 different IPv4 originator addresses* and can thus distribute the scan traffic. This allows the search to stay below the radar of many intrusion prevention systems and thus increases the scan data quality significantly.

As soon as a valid packet arrives from a destination address, the application scanner is started. The application scanner supports more than 60 different protocols and establishes full application connections with the goal of reading as much identification information as possible from the system. Most of the protocols we implemented ourselves; for several standard protocols, we used the Zgrab implementation [21].

After the IPv4 addresses of a work unit are processed, the scan results are sent to the aggregator. The aggregator collects all results from search nodes and checks them for consistency. Then the data is enriched with other information from open sources in JSON format.

For example, the IPv4 scan data is enriched with the INETNUM and WHOIS information from the RIRs (RIPE, ARIN, AfriNic, etc). Possible inconsistencies within the databases, such as overlapping prefixes, are resolved according to a self-developed method defined in [12]. In addition, the BGP data valid at the respective time is stored for each IPv4 address. This includes the BGP prefix annotated at the time, including all available autonomous system data. As a data source for the BGP information, the data of the Cooperative Association for Internet Data Analysis (CAIDA) [10] is merged and processed. In addition, reverse DNS records and Geo-IP information are added to each discovered active IPv4 address of the respective scan. All data is stored in a NOSQL-based ElasticSearch database, which can be duplicated as an on-premises solution for private discretionary analysis – commonly needed by, for instance, defense organizations – at any time.

For the analysis of the hospital data, a separate subfrontend called Inspector was developed, to make this complex task easier for our human analysts. The Inspector only receives the names of the hospitals and the respective domain name as input. Subsequently, all relevant entries in the database, such as the WHOIS description field, or the common names of the collected TLS certificate information or the atomic system data, are analyzed for membership of the respective target set using self-developed advanced Big Data algorithms. In parallel, all subdomains of the added domains are searched. This is done by special best guess algorithms or by searching known public certificate databases. The Inspector had to be created, as our analysts had to take a huge list of potential healthcare target organizations into account.

The final step is about vulnerability detection: after all network services of the defined reconnaissance targets, in our case hospitals and other healthcare providers, had been identified, the system descriptions or version strings read out were compared with the National Vulnerability Database (NVD) of NIST [19]. Through this step, all potential known vulnerabilities in detected software systems are identified.

4. METHODOLOGY OF ATTACK SURFACE DETECTION OF HOSPITALS AND CLINICAL PROVIDERS

To identify the attack surface of the German hospitals, the German hospitals themselves first had to be identified. Therefore, we chose as a starting point the German DIVI registry [9], which was first established during the COVID-19 pandemic.

The DIVI Intensive Care Register records the free and occupied treatment capacities in intensive care medicine at about 1,300 acute hospitals in Germany on a daily basis. During the pandemic and beyond, the registry makes it possible to identify bottlenecks in intensive medical care, regionally and/or temporally. Thus the DIVI Intensive Care Registry creates a valuable basis for response and data-driven action control in near real-time since April 2020.

Our approach was the following: we extracted over 1,300 names of German hospitals with COVID-19 intensive care units from the DIVI Register. We then manually searched for the main website or domain of the corresponding hospital names and added them to the DIVI Registry data.

In the next processing step, the names and domain information were entered into the Inspector. The Inspector then analyzed a total of 89 different port/protocol scans. A sizable amount of data – 1,483 GB – was analyzed on a system with 1 TB Ram, 64 CPU cores, 40 TB SSD storage and 72 TB HDD storage. The total computing time of the whole system was about 16 hours.

Table I is a listing of the port/protocol combinations for which global scans for about 2.8 billion routed IPv4 addresses have been conducted.

After identification of the associated network services based on the certificate information and WHOIS and BGP/AS data, as well as the extended detection of subdomains, it was possible to collect additional information about other hospitals. For example, the cryptographic TLS certificate of Hospital A might also include the domain of another Hospital B of the same provider. Furthermore, generic search terms such as “hospital” and “clinic” were added. In addition, the results were manually searched, and any false positives were eliminated. This approach made it possible to extend the analysis of 1,300 hospitals of the DIVI registry to 1,555 hospitals.

TABLE I: SCANNED TCP AND UDP PORTS DURING ATTACK SURFACE MAPPING

http-1000
http-5985
http-7547
http-80
http-8008
http-8080
http-8088
http-8888
https-1443
https-443
https-4433
https-4434
https-4444
https-5986
https-8443
dnstcp-53
elastic-9200
eniptcp-44818
fox-1911
ftp-21
openport-873
openport-9200
openport-587
bacnet-47808
bigip-443
cve20205902-443
dnp3-20000
imap-143
ipmi-623
ipp-631
kibana-5601
knx-3671
ldap-389
ldapudp-389
modbus-502
mongodb-27017
mssqludp-1433
mssqludp-1434
mysql-3306
netbios-137
ntp-123
oracledb-1521
pop3-110
openport-6379
openport-8009
postgres-5432
qnapvuln-8080
redis-6379
s7-102
samba-445
snmpv1-161
snmpv2-161
ssh-22
ssh-22022
ssh-2222
sworionrest-17778
telnet-23
telnet-2323
telnet-4786
telnet-5938
telnet-7070
upnp-1900
winrm-5984
openport-993
openport-995
openport-5900
openport-5984
openport-1025
openport-111
openport-11211
openport-11711
openport-1201
openport-135
openport-139
openport-1433
openport-1521
openport-1720
openport-1723
openport-199
openport-2012
openport-27017
openport-3306
openport-3389
openport-445
openport-469
openport-5037
openport-5432
openport-5555
openport-5601

5. DATA SECTION

Our analysis of the 1,555 German hospitals revealed a digital attack surface of 13,497 network services, or 8.7 network services per hospital on average. Figure 2 shows the distribution of the main service banner groups of all identified hospital network services which were identified by executing a full handshake.

FIGURE 2: DISTRIBUTION OF THE MOST COMMON DETECTED SERVICE BANNER GROUPED BY MAJOR SERVICE APPLICATION

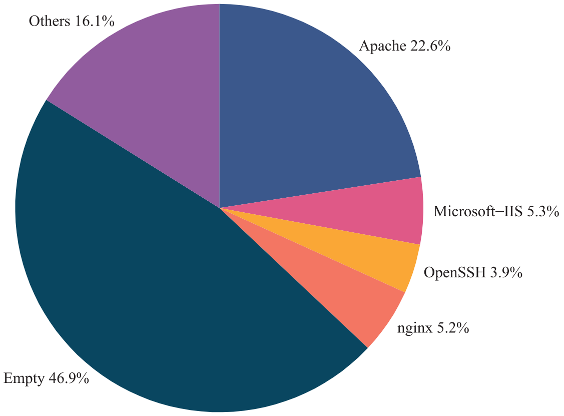

Approximately 47 percent of all collected service banners are empty and thus comply with the common best-practice approach of not disclosing any software version information via service banner. This approach is very important because it makes it more difficult for attackers to identify the software used. This makes it subsequently harder for a potential attacker to determine the proper exploit/malware to use in an attack attempt. This is especially true for the use of automated attack scripts, often used by automated botnets.

We identified 1,228 hospitals and hospital operating companies that had network services that could be directly located. Approximately 300 other hospitals had no network services of their own but only those that could be assigned to joint operating companies. However, since we do not know how the networks of the joint hospital operating companies are related to the hospitals, we consider the whole operating company as a single hospital. Thus we technically analyze 1,228 hospital entities and operating companies representing up to 1,555 different hospitals. Of the 1,228 hospitals, 447 had vulnerable network services. This means that 36.4 percent of all identified hospitals and hospital operating companies have vulnerabilities.

Figures 3, 4, and 5 show the version distribution of the three most common web servers: Apache httpd, Microsoft IIS, and nginx. A well-known problem in the industrial and (to a certain degree) the healthcare sector became visible quite early in our analysis: outdated services for which end-of-support had already been announced. The most noteworthy candidates we identified included Apache httpd Version 2.2.x, which became end-of-support in December 2017, or Microsoft Internet Information

Services 6.0, which became end-of-support in June 2015. It is unclear, however, why we found those legacy services on Internet-facing systems, as the issue of patch and update difficulty typically affects mainly internal medical components, not Internet infrastructure.

FIGURE 3: VERSION DISTRIBUTION OF DETECTED APACHE WEB SERVERS, WITH ROUGHLY ONE-THIRD HAVING KNOWN VULNERABILITIES. NOTE THAT 2,092 APACHE SERVERS (68.43 PERCENT) RESULTED IN AN UNDEFINED VERSION AND ARE NOT INCLUDED

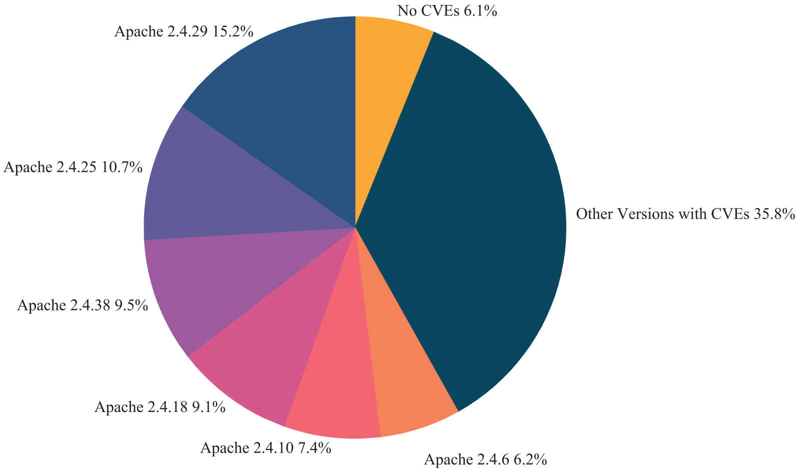

FIGURE 4: DISTRIBUTION OF DETECTED VERSIONS OF MICROSOFT INTERNET INFORMATION SERVICES (IIS) WEBSERVER, INDICATING CURRENT AS WELL AS END-OF-SUPPORT VERSIONS IN OPERATION

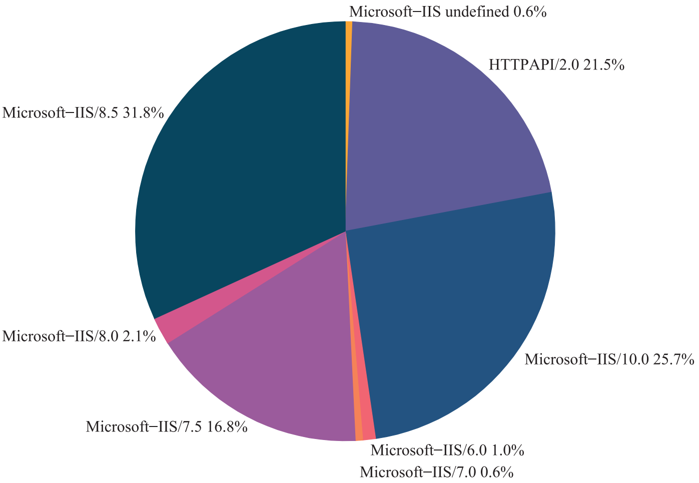

FIGURE 5: DISTRIBUTION OF DETECTED VERSIONS OF NGINX WEBSERVER, INDICATING CURRENT AS WELL AS END-OF-SUPPORT VERSIONS IN OPERATION. NOTE THAT 444 NGINX SERVERS (62.62 PERCENT) RESULTED IN AN UNDEFINED VERSION AND ARE NOT INCLUDED.

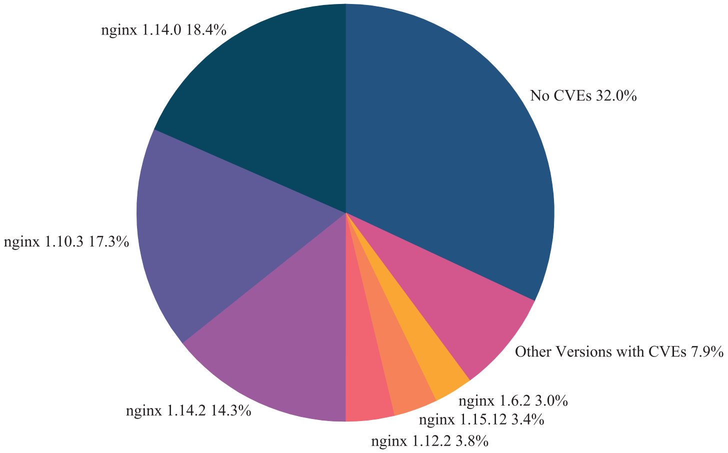

Figure 6 shows the geographic location of all 1,300 hospitals of the DIVI register. This clearly shows that there is a high density of hospitals, particularly in the densely populated regions of western Germany and in the German metropolitan areas of Hamburg, Berlin, and Munich (see Figure 6a). The image on the right (Figure 6b) shows the DIVI registry hospitals with vulnerabilities on the map. It is easy to recognize that hospitals in both metropolitan areas and rural areas are affected.

FIGURE 6: GEOLOCATION OF HOSPITALS, NETWORK SERVICES, AND VULNERABILITIES

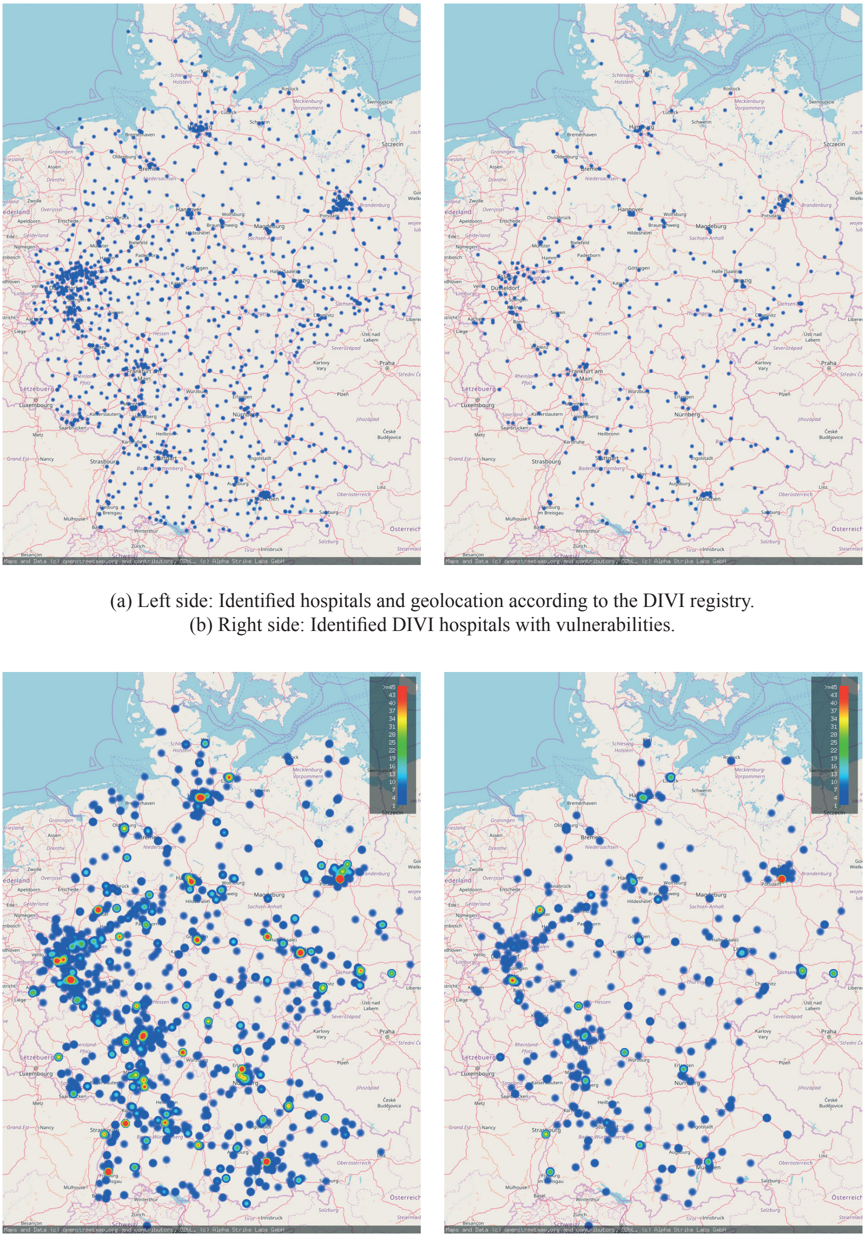

(a) Left side: Identified hospitals and geolocation according to the DIVI registry.
(b) Right side: Identified DIVI hospitals with vulnerabilities.

(c) Left side: All network services identified and the approximate Geo-IP location as heatmap.
(d) Right side: Geographical location of the network services with vulnerabilities as heatmap.

In contrast to Figure 6a and 6b, Figure 6c and 6d represent an overview of all 1,555 identified hospitals and their 13,597 network services, which were assigned a geo-coordinate via a Geo-IP resolution. For the Geo-IP resolution, the commercial version of the Maxmind DB [16] with increased resolution was used. Figure 6c shows the network services of all hospitals analogous to Figure 6a, whereas Figure 6d shows only the network services with vulnerabilities.

The main difference between Figure 6a and 6d is that Figure 6a only shows the hospitals of the DIVI Registry and their geographical location. Figure 6d, however, shows a heat map of all identified or vulnerable network services of German hospitals. A comparison of the two graphs clearly reveals that the distribution in the heat map is somewhat smaller, but both graphs show that both rural regions and metropolitan areas have hospitals with vulnerabilities.

First, the following overall CVSS vulnerability statistics should be noted:

TABLE II: CVSS DISTRIBUTION OVERVIEW

CVSS SCORE	Number of vulnerable services
9.0–10 (critical) 7.0–8.9 (high) 4.0–6.9 (medium)	931 443 518
Total vulnerable services:	1,892

Our analysis yielded a total of 1,892 vulnerable services, with nearly half of the vulnerable services carrying a CVSS score of 9 or 10, thereby potentially containing critical vulnerabilities, depending on their version number.

Next, we explore if there is any correlation between the number of identified CVEs (Common Vulnerabilities and Exposures, a reference method for publicly known information-security vulnerabilities and exposures) and the size of the clinical institution. The allocation of the number of beds was taken from the German Hospital Register [8]. An examination of the hospitals with vulnerabilities in relation to their bed capacity shows that hospitals with vulnerabilities represent a total of about 167,000 beds. This represents 32 percent of the approx. 520,000 available hospital beds in Germany. Figure 7 shows the number of identified CVEs in relation to the size of the respective hospitals based on the number of beds. For a better visualization, only hospitals with up to 1,800 beds are drawn; there are only a few facilities with more beds.

Since, naturally, there are more smaller hospitals, there are more data points in the left-hand area of the figure. For better visibility, a detailed representation of this area is shown in Figure 7b.

FIGURE 7: NUMBER OF VULNERABILITIES IN IT SYSTEMS IN HOSPITALS

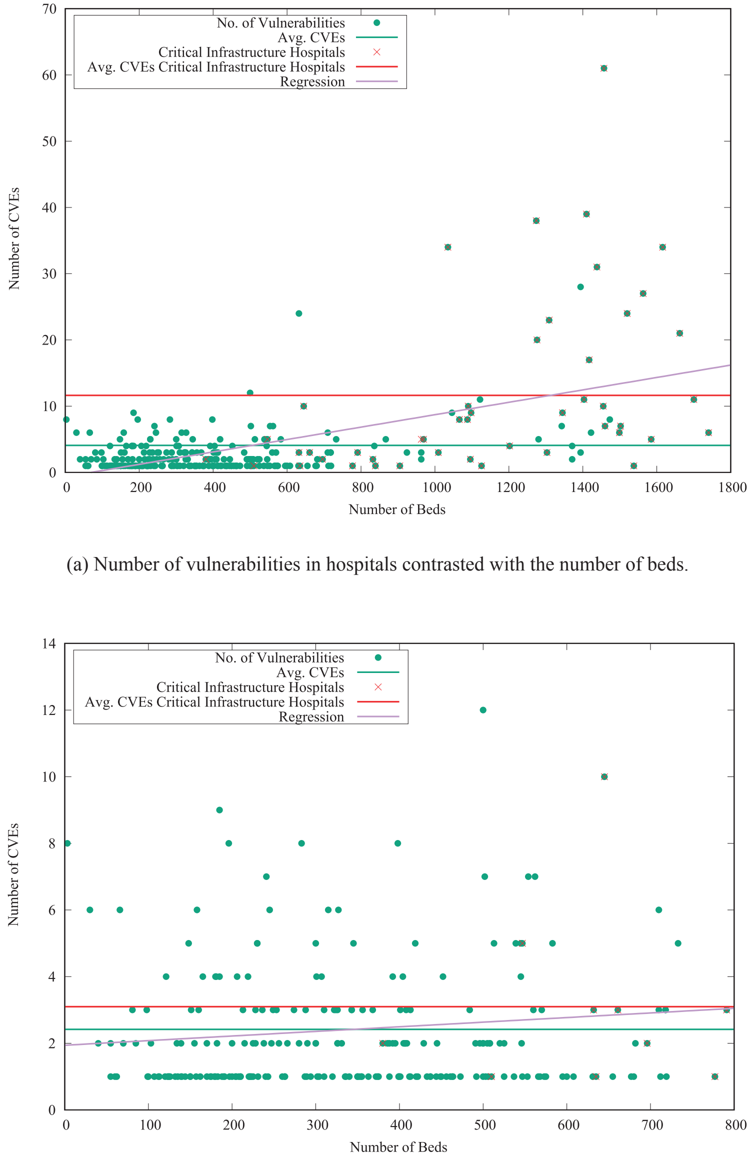

(a) Number of vulnerabilities in hospitals contrasted with the number of beds.

(b) Detail view of the number of vulnerabilities in hospitals with up to 800 beds.

A first look at the data initially reveals an unsurprising trend: As the number of beds increases, so does the number of vulnerabilities found in the IT environments of the respective hospitals. This can probably be explained by the fact that larger hospitals with more beds also typically have more specialized medical departments and corresponding IT equipment, which thus increases the number of IT devices and services and thus the potential attack surface. The regression line of this increase is drawn in the figures correspondingly.

But if we now look at the detailed view in Figure 7b, we notice that a corresponding increase in vulnerabilities in IT systems is much lower among hospitals with up to 800 beds.

Here hospitals in all size ranges have varying numbers of vulnerabilities, without any discernible pattern. We might infer from this that at smaller hospitals, the number of existing vulnerabilities is more likely to depend on the quality of the respective IT service providers, or on specific software products.

With respect to the significantly increasing numbers of vulnerabilities at large hospitals, however, especially at those with more than 1,000 beds, it is apparent that these affect university hospitals in particular. This suggests that the higher CVE figures also reflect the need for more systems and, especially in the research sector, more diversified IT systems and customized or less commonly available software.

With respect to German legislation, the data in Figure 7a offers yet another perspective for analysis: due to the special need for protection of the basic services necessary for modern society, such as electricity water supply, telecommunications, and healthcare, the Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik, BSI) criticalness regulation (KRITIS Act [7]) defines facilities in Germany that are obligated to implement minimum standards and measures in accordance with the BSI Act [2] to ensure sufficient IT security. In the healthcare sector, facilities with more than 30,000 in-patient cases per year are considered critical infrastructure.

Thus an interesting question arose: are higher liabilities resulting from the KRITIS Act reflected in a lower visible CVE attack surface?

In order to evaluate this, the facilities that belong to KRITIS based on the number of cases according to the German Hospital Register [8] were marked accordingly in Figure 7a. Of course, large facilities such as university hospitals fall into this category, but so do some other, smaller facilities. Surprisingly, while the aforementioned accumulation of vulnerabilities can be seen at university hospitals, smaller institutions also feature systems with more vulnerabilities than average.

For example, an average of 11.63 CVEs was identified for hospitals up to 1,800 beds belonging to KRITIS, while the average value for all of the hospitals up to 1,800 beds analyzed is 4.08. A similar picture emerges when looking at the detail section of smaller hospitals in Figure 7b. While the average number of CVEs at the KRITIS hospitals is 3.1, all analyzed hospitals with up to 800 beds have an average of 2.42 CVEs.

6. DISCUSSION OF RESULTS AND CONCLUSION

In our results section, we first want to acknowledge the known limitations and constraints of our analysis, beginning with the number of vulnerable services our DCS identified (1,892). Firstly, it must be noted that vulnerability identification is done fully automatically through service and banner mapping and CVE entry. In cases where patches have been backported or the administrator has arbitrarily changed banner information, the given CVE match indication naturally would not reflect the actual vulnerability state. For cases like this, we coined the term "Schrödinger vulnerability," which is explained later in this section. Therefore, it may be assumed that the overall number might be a bit lower due to backports or banner changes. On the other hand, other attack vectors such as misconfiguration of services or the use of weak passwords, which are still regularly found today and can represent a high risk for an IT system, are not included in our research. Therefore, taking into account these considerations, our results can also be seen as a *lower bound* for the vulnerabilities of the systems, and their effective exposure can be even higher. Secondly, although DCS uses a number of very well-proven port and service identification methods, in cases where fingerprinting fails, this may create a situation where vulnerability identification is not always accurate.

While we have only analyzed IPv4 addresses in the present work, we are working on the implementation for scanning procedures for the IPv6 protocol. Since its address range is no longer (even approximately) completely scannable due to its sheer size, new and different scanning strategies are required here to reduce and optimize the search space. Efforts in this direction are already underway.

However, regarding the results, it is also important to recognize the great advantage of our method, which is typically unproblematic from a legal point of view due to the evaluation of information provided publicly by Internet-facing services only. By contrast, even if actors such as intelligence services of foreign countries are probably not bothered by this limitation in various cases, performing vulnerability scans to explicitly search for and thereby trigger vulnerabilities without the permission of the owner of the respective service is problematic for us.

In our research using DCS, we would like to coin the term “Schrödinger vulnerability.” In quantum mechanics, “Schrödinger’s cat” is a thought experiment conceived by the Austrian-Irish physicist Erwin Schrödinger in 1935 in a discussion with Albert Einstein [26]. In the thought experiment, a hypothetical cat can be considered simultaneously alive and dead because it is associated with a random subatomic event that may or may not occur.

In the thought experiment, a closed box contains a cat and an unstable atomic nucleus, which decays with a certain probability within a certain period of time. The decay would trigger the release of poison gas by means of a Geiger counter and thus kill the cat.

The thought experiment is based on the fact that whenever a system can assume two different states, the coherent superposition of the two states then also constitutes a possible state. It is therefore only an actual observation or measurement being conducted that can distinguish the two original states, as the system may assume either one.

Analogous to this thought experiment, we now consider an IT system with a network service (aka the cat) that has a vulnerability according to its transmitted version in the service banner. With that, however, we cannot know whether the administrator (aka the atomic core) of the IT system has provided the network service with a security patch, because many security updates do not update the communicated software version in the service banner. Consequently, the IT system is in a state of superposition, as it is both vulnerable and non-vulnerable at the same time. Only when the system is subjected to a thorough audit – for example, by analyzing the exact version level or patch level – can we distinguish between the original states, and until then, the system may adopt either a state of vulnerability or one of non-vulnerability.

Consequently, we have to call all vulnerabilities, which in the context of this analysis were mostly identified via the service banner, Schrödinger vulnerabilities or potential vulnerabilities, as they put a system in a state of vulnerability and non-vulnerability simultaneously. From the authors’ perspective, all Schrödinger vulnerabilities should therefore, until further audited, be considered a cyber risk.

Further distinctions could be made for future work. For example, evaluations could be created to show the proportion of systems that can no longer be supplied with current software versions or security updates because the software products have already been marked as end-of-support by the manufacturer.

Let us now summarize the results and findings. First of all, when looking at the attack surface of German clinical providers, the analysis reveals many vulnerabilities and quite high CVSS ratings. Looking at the most noteworthy system occurrences from a security point of view reveals, for instance, two Windows XP operating systems (CVSS 10.0/End of Support since 2015!), open Jitsi VideoChat servers (CVSS 6.11), open unauthenticated squid proxies (CVSS 10.0) allowing proxy misuse, outdated Apache and PHP configurations (CVSS 9.8), direct accessible Intelligent Platform Management Interface (IPMI) login pages, Citrix XenAPP remote access (CVSS 10.0), and direct web links to RDP connections (CVSS 9.8), to give just a few concrete examples.

The main designated and essential function of clinical institutions is healthcare and not IT security. However, the data seem to indicate that there is still a need for better attack surface and vulnerability management, as **approximately 32 percent of the analyzed services were determined to be vulnerable to various degrees, and 36 percent of all hospitals showed vulnerabilities.**

As mentioned, through our analysis we can confirm that healthcare institutions are also affected to a certain extent by the issue of legacy services for which end-of-support was announced years ago and for which security updates are therefore no longer provided, increasing security risk.

Unsurprisingly, larger institutions have more IT systems, potentially leading to a larger attack surface; this was clearly visible in our analysis as well.

Finally, a rather interesting result of our analysis was the fact that hospitals belonging to German critical infrastructure, indicated through their assignment from the KRITIS Act, had **notably higher-than-average vulnerability figures, based on CVE numbers**, among the hospitals we analyzed. We found this result striking, as we had assumed that KRITIS hospitals and clinics would have a much better IT security posture, resulting in lower average CVE numbers than the other hospitals, as they are designated as being critical.

Our analysis concludes that even in 2020, despite its increased criticalness and increased regulation efforts, the German healthcare sector, unfortunately, presented and contained a certain visible amount of attack surface. This attack surface may translate into a national security risk if abused systematically by an intelligent adversary. It is therefore advisable from a state-level risk management perspective to regularly conduct reconnaissance in cyberspace on all organizations that have been determined critical or essential for a nation.

REFERENCES

[1] Cyber and Infrastructure Agency (CISA), "CISA: Alert (AA20-302A) Ransomware Activity Targeting the Healthcare and Public Health Sector." Accessed: Nov. 15, 2020. [Online]. Available: https://us-cert.cisa.gov/ncas/alerts/aa20-302a

[2] *Act on the Federal Office for Information Technology (BSI Act – BSIG)*. Accessed: Jan. 6, 2021. [Online]. Available: https://www.bsi.bund.de/EN/TheBSI/BSIAct/bsiact_node.html

[3] Jim Bates, "Trojan Horse: AIDS Information Introductory Diskette Version 2.0," *Virus Bulletin*, vol. 6, pp. 1143–1148 (1990).

[4] Coalition Inc., "Cyber Insurance Claims Report H1 2020." Accessed: Feb. 26, 2021. [Online]. Available: https://info.coalitioninc.com/rs/566-KWJ-784/images/DLC-2020-09-Coalition-Cyber-Insurance-Claims-Report-2020.pdf

[5] "Common Vulnerability Scoring System SIG." Accessed: Jan. 6, 2021. [Online]. Available: https://www.first.org/cvss

[6] Andrada Fiscutean, "A history of ransomware: motives and methods behind evolving attacks," *CIO East Africa*, July 28, 2020. Accessed: Jan. 6, 2021. [Online]. Available: https://www.cio.co.ke/a-history-of-ransomware-the-motives-and-methods-behind-these-evolving-attacks

[7] Deutsche Krankenhaus TrustCenter und Informationsverarbeitung GmbH (DKTIG), "BSI-Kritisverordnung (BSI-KritisV)." Accessed: Jan. 4, 2021. [Online]. Available: https://www.gesetzeim-internet.de/bsi-kritisv

[8] Deutsche Krankenhaus TrustCenter und Informationsverarbeitung GmbH (DKTIG), "Deutsches Krankenhaus Verzeichnis." Accessed: Jan. 3, 2021. [Online]. Available: https://www.deutscheskrankenhaus-verzeichnis.de

[9] Robert Koch Institut, "DIVI Intensivregister." Accessed: Jan. 4, 2021. [Online]. Available: https://www.intensivregister.de/#/index

[10] "Cooperative Association for Internet Data Analysis (CAIDA)." Accessed Jan. 3, 2021. [Online]. Available: https://www.caida.org

[11] Jessica Davis, "The 10 biggest healthcare data breaches of 2020," *Health IT Security News*, Dec. 10, 2020. Accessed: Jan. 4, 2021. [Online]. Available: https://healthitsecurity.com/news/the-10-biggest-healthcare-data-breaches-of-2020

[12] Johannes Klick, "Towards Better Internet Citizenship: Reducing the Footprint of Internet-wide Scans by Topology Aware Prefix Selection," Proceedings of the 2016 Internet Measurement Conference, pp. 421–427 (2016).

[13] Forescout Research Labs, "Connected Medical Device Security: A Deep Dive into Healthcare Networks," Tech. Rep. Forescout Technologies, Inc., 2020. Accessed: Dec. 23, 2020. [Online]. Available: https://www.forescout.com/company/resources/connectedmedical-device-security-a-deep-dive-into-healthcare-networks

[14] Joanne Lim, "An Engineering Disaster: Therac-25," 1998. Accessed: Apr. 7, 2021. [Online]. Available: https://www.bowdoin.edu/~allen/courses/cs260/readings/therac.pdf

[15] GLEG Ltd., "Gleg Security @GlegExploitPack." Accessed: Jan. 6, 2021. [Online]. Available: https://twitter.com/GlegExploitPack

[16] Maxmind, "GeoIP Databases and Services." Accessed: Jan. 4, 2021. [Online]. Available: https://www.maxmind.com/en/geoip2-services-and-databases

[17] Melissa Hanson, "Swansea Police Department pays ransom to computer hackers," *Boston Globe*, Nov. 19, 2013. Accessed: Jan. 3, 2021. [Online]. Available: https://www.bostonglobe.com/metro/2013/11/19/swansea-police-pay-ransom-open-files-locked-hackers/7bOdi8i7foNkTmdnokMAkP/story.html

[18] Mitchell Clarke and Tom Hall, "It's not FINished: The Evolving Maturity in Ransomware Operations." Accessed: Jan. 4, 2021. [Online]. Available: https://www.blackhat.com/eu-20/briefings/schedule/index.html#its-not-finished-theevolving-maturity-in-ransomware-operations-21500

[19] "National Vulnerability Database." Accessed: Jan 6, 2021. [Online]. Available: https://nvd.nist.gov

[20] Lily Hay Newman, "Ransomware hits dozens of hospitals in an unprecedented wave," *Wired*, October 29, 2020. Accessed: Nov. 17, 2020. [Online]. Available: https://www.wired.com/story/ransomware-hospitals-ryuk-trickbot/

[21] ZMAP Project. "ZGrab Repository." Accessed: Jan 4, 2021. [Online]. Available: https://github.com/zmap/zgrab2

[22] William Ralston, "The untold story of a cyberattack, a hospital and a dying woman," *Wired*, Nov. 11, 2020. Accessed: Nov. 18, 2020. [Online]. Available: https://www.wired.co.uk/article/ransomware-hospital-death-germany

[23] Kim Zetter, "Why hospitals are the perfect targets for ransomware," *Wired*, March 30, 2016. Accessed: Nov. 16, 2020. [Online]. Available: https://www.wired.com/2016/03/ransomware-why-hospitals-are-the-perfect-targets/

[24] U.S. Department of the Treasury, Office of Foreign Assets Control (OFAC), "Advisory on Potential Sanctions Risks for Facilitating Ransomware Payments." Accessed: Jan. 3, 2021. [Online]. Available: https://home.treasury.gov/system/files/126/ofac_ransomware_advisory_10012020_1.pdf

[25] Phil Goodwin and Andrew Smith, "The State of IT Resilience," White Paper, International Data Corporation (IDC), August 2019. Accessed: Feb. 26, 2021. [Online]. Available: https://www.zerto.com/wp-content/uploads/2019/07/State_of_IT_Resilience_2019.pdf

[26] E. Schrödinger, "Die gegenwärtige Situation in der Quantenmechanik," *Naturwissenschaften*, vol. 23, pp. 807–812 (1935) (in German). Accessed: Mar. 8, 2021. [Online]. Available: https://doi.org/10.1007/BF01491891