2020 RESEARCH REPORT

Connected Medical Device Security:

A Deep Dive into Healthcare Networks

By Forescout Research Labs

FORESCOUT
Active Defense for the Enterprise of Things™

CONTENTS

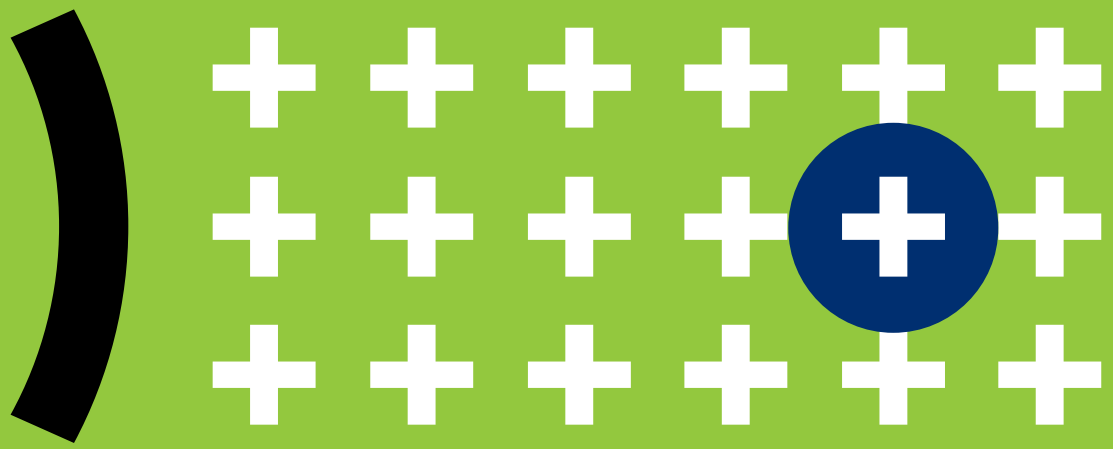

CONTENTS

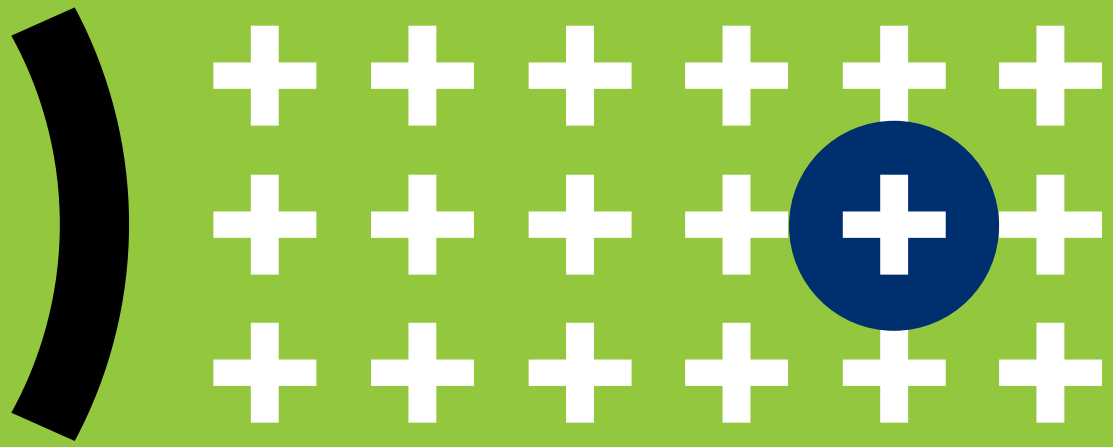

EXECUTIVE SUMMARY

Healthcare delivery organizations (HDOs), such as hospitals and clinics, are complex organizations where a **broad range** of Information Technology (IT), Internet of medical things (IoMT), Operational Technology (OT) and Internet of Things (IoT) **devices are increasingly interconnected** [1] [2].

The growing number and diversity of devices in HDOs have introduced new **cybersecurity risks** [3] [4] [5]. The ability to **compromise devices and networks** and the possibility of **monetizing patient data** [6] [7] have led to an **increase in the number and sophistication of cyberattacks targeting healthcare delivery organizations** in recent years [8]. As a result, 82% of U.S. hospitals report having a significant security incident in 2018 or 2019 [9].

Changes in HDO networks in 12 months

In April 2019, Forescout Research Labs analyzed the security of healthcare delivery organizations using the Forescout Device Cloud and **found major risks** associated with the use of **legacy systems** and **insufficient segmentation** [10].

One year later, we **applied a similar analysis to the most recent data in our Device Cloud**, which led to the findings in this report of some overall **improvements in patching and network segmentation**. However, we still saw **many examples of poorly segmented networks** with a mix of personal and sensitive healthcare devices, including devices with default passwords, which is a top IoT cyber risk [11].

Given these results, we **decided to closely analyze network traffic patterns in several large HDOs** to better understand how lack of segmentation coupled with observed issues such as the use of **insecure protocols** and inappropriate **external communications** leads to increased cyber risk, an enlarged attack surface and difficult-to-secure networks.

The **key findings** of this report are the following:

1. Most healthcare networks have upgraded to **Windows 10 over the past year** and embraced some segmentation with **the number of VLANs increasing** when compared to 2019.
2. There are still many examples of **network segmentation** issues, including a mix of personal and medical devices in healthcare segments.
3. The **analyzed healthcare delivery networks heavily used insecure protocols for both medical and non-medical** network communications. We also found examples of **sensitive external communication**.
4. Based on the previous findings, we **demonstrate some easy-to-accomplish attacks targeting point-of-care testing devices and patient monitors, some of the most commonly used IoMT devices in an HDO**. Although similar issues have been demonstrated for a few well-known protocols, we extend this to lesser-known protocols that interconnect a multitude of devices.

We conclude by discussing **effective strategies to reduce cybersecurity risk and defend healthcare networks** from cyberattacks.

1. Architecture of a typical healthcare network

The typical healthcare delivery organization observed in this study is a **hospital inpatient facility with on-site labs and specialty units such as ICUs, burn units and others**. These healthcare delivery organizations contain an **average of 20,000 devices**, including IT, IoMT, OT, and IoT devices. **IT devices** exchange **highly sensitive data** (e.g., patient health records and financial information), whereas **IoMT, OT and IoT** devices are used for **diverse functions** such as building automation, guest entertainment, patient monitoring, and healthcare delivery.

In this section, we explore the devices we **typically found at the sites in our study** (Section 1.1) and some of the **potential threats** to these networks (Section 1.2).

1.1. Observed devices

Figure 1 shows a sample of devices typically found in the studied HDO networks.

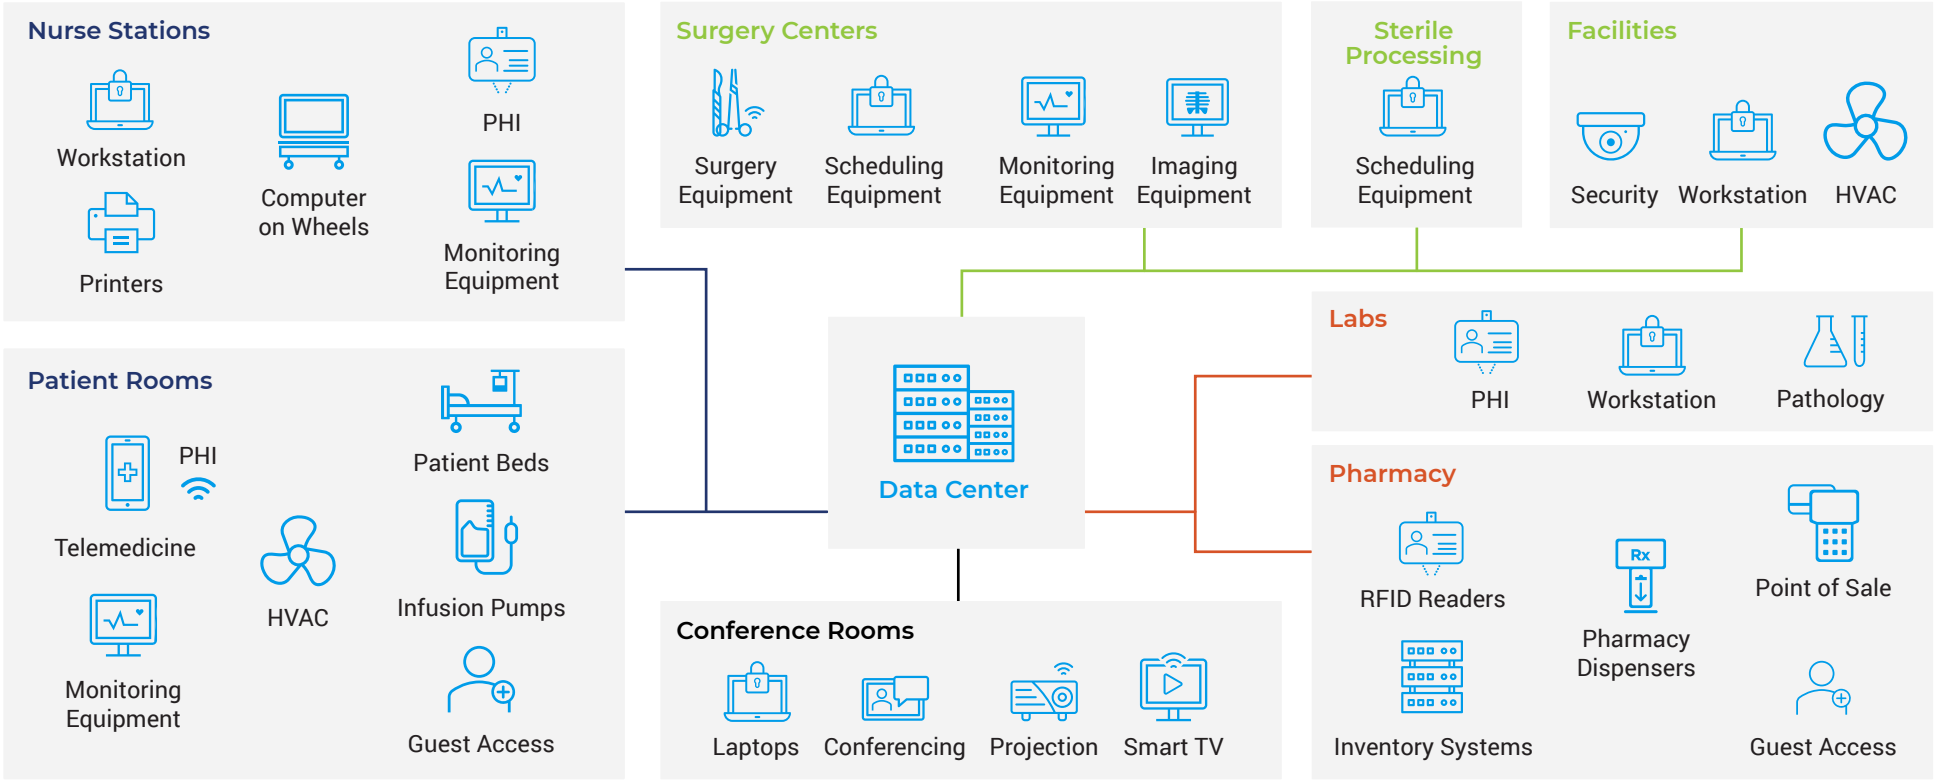

Figure 1 – Devices in a typical healthcare network.

Right in the center is the **data center**, where Electronic Health Records (EHR) and Electronic Medical Records (EMR) systems are present and contain the **crown jewels in healthcare**. From a threat perspective, that's the primary data store. That's where the Personal Health Information (PHI), payment information and all the data required to generate a superbill are stored.

Then there are **connected medical devices** in patient rooms, nurse stations, surgery centers, pharmacies, labs, and countless other locations. These devices may **support clinical care**, such as insulin pumps, heart defibrillators, ventilators, and any equipment saving or sustaining life, **or gather and monitor patient information** (such as vital signs and test results) to alert and inform clinical staff. These include patient monitors, laboratory equipment, imaging devices, and more.

However, many IoT devices on healthcare delivery organization networks **aren't medical devices at all**, including cafeteria and **pharmacy point-of-sale systems, vending machines, ATMs, gift shop kiosks, an array of physical security devices, and smart building systems** for energy and power management, HVAC, and backup generators. For more details on how some of these smart building systems increase the attack surface of organizations, see our previous research reports [12] [13].

1.2. Network diagram and potential threats

Figure 2 shows a **simplified network diagram** of a typical HDO network and some of the **potential threats** faced by security teams. Notice that the network has an **IT section** – containing traditional devices such as PCs, mail and web servers – and a **Clinical/ IoT section**, containing medical devices as well as traditional workstations and devices such as IP cameras, wireless routers and printers.

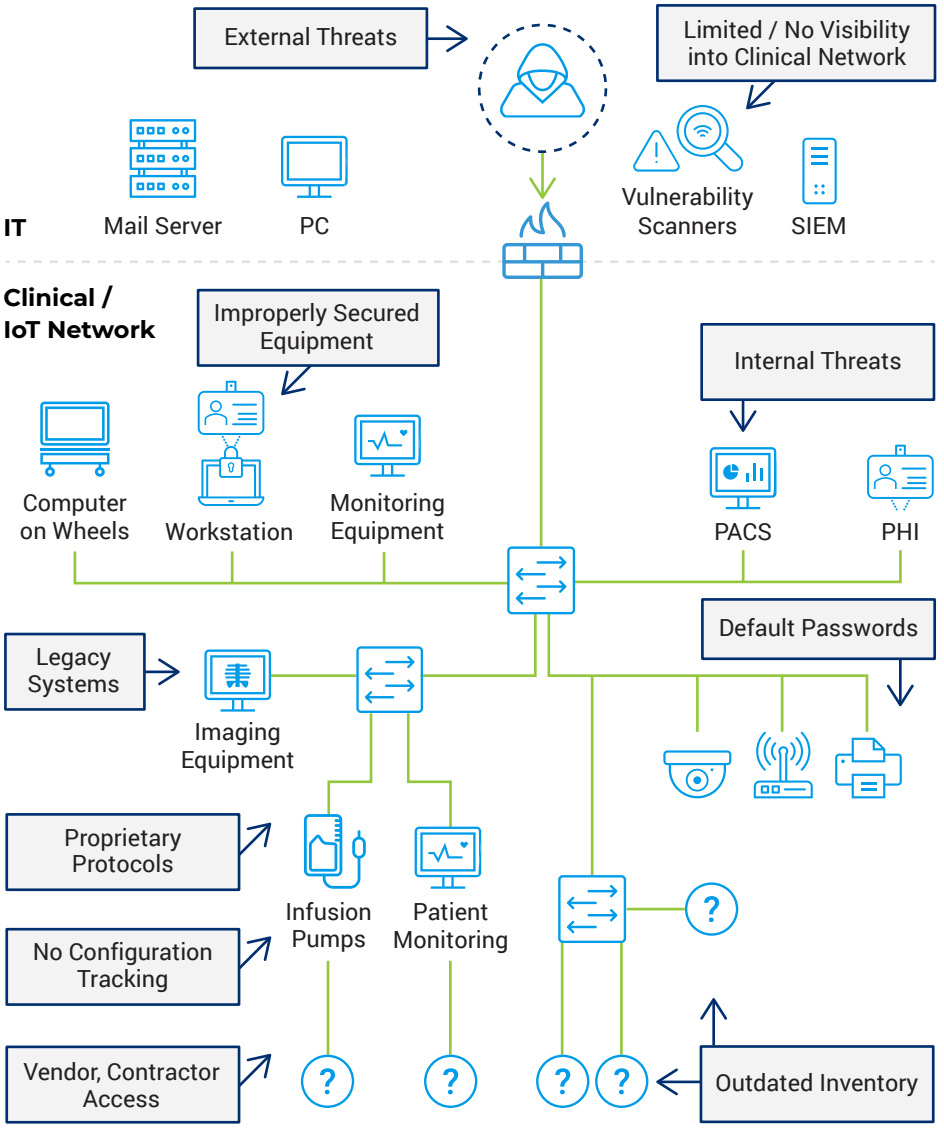

Figure 2 – Network diagram with potential threats.

Some potential threats that are shown in the Figure (in purple boxes) include:

- **External Threats –** Many malicious actors have motivations to attack healthcare delivery organizations [9] [14]. Individual cybercriminals or criminal organizations usually try to reap money from cyberattacks, either directly via ransomware and cryptomining or indirectly by selling stolen information or access to infected computers with botnets.
- **Internal Threats, Vendor and Contractor Access –** It's important to note that when it comes to attack surfaces, HDOs aren't just running emergency rooms and – medical clinics – they're running remote access VPNs for vendor support and extensive back offices for administrative functions that require common (yet privileged) workstations. These open the network to the possibility of attacks by internal bad actors, which may also have a financial motivation, but may also have other goals such as sabotage.
- **Legacy Systems, Improperly Secured Equipment, and Default Passwords –** These are some of the most important technical issues faced by security personnel in HDOs. First, many medical devices are legacy systems that cannot be patched due to availability or certification requirements (see more details in Section 2.1.1). These systems tend to be easy targets for attackers because of well-known vulnerabilities that cannot be patched and the existence of commoditized exploits for them. Second, even systems that can be patched are often improperly secured, such as unmanaged endpoints, or improperly configured, such as devices with known default passwords that can be obtained online [15] and devices exposing network services that are not required.

Defending against these threats requires the **use of multiple security tools**, such as **network monitoring, vulnerability scanners and SIEM systems**, as shown in the top right corner of Figure 2. However, these tools face **difficulties in the healthcare world** because of issues such as **proprietary protocols, lack of configuration tracking, outdated inventory and limited visibility into the clinical network**.

The results are a **lack of visibility into specific information about medical devices** (specific models, software versions, and serial numbers, etc.) and their network activity, which **makes it difficult to detect vulnerabilities, segmentation problems and even ongoing attacks**.

2. Analyzing healthcare devices and networks

In this section, we analyze healthcare networks using two data sources: the Forescout Device Cloud, containing data for about 3.3 million devices in hundreds of healthcare networks (Sections 2.1 and 2.2), and a detailed analysis of network traffic from several large healthcare delivery organizations (Section 2.3).

2.1. Device cloud findings: improvements from 2019

In April 2019, we analyzed a subset of our Device Cloud containing **75 healthcare deployments**, including **over 10,000 Virtual Local Area Networks** (VLANs) and **1.5 million devices**. In short, the analysis found major risks associated with the use of legacy systems and insufficient segmentation [10].

One year after the initial study, we performed a **similar analysis using recent data in our Device Cloud** and compared the results with what was seen last year in terms of versions of the Windows OS (Section 2.1.1) and the number of network segments (Section 2.1.2).

2.1.1. Windows versions

When we analyzed the data in 2019, Microsoft had announced the end of support for Windows 7, Windows 2008 and Windows Mobile for January 14, 2020 [16]. At that time, we saw 71% of Windows devices in our sample of healthcare networks running soon-to-be unsupported versions.

This year, after the end-of-support date passed, we analyzed the effect this had on healthcare deployments. **Figure 3** shows the number of devices running Windows OS versions that are **still supported ("The Good," e.g., Windows 10), supported only via the** paid Extended Security Update [17] **program ("The Bad," e.g., Windows 7)** or **totally unsupported ("The Ugly," e.g., Windows XP)** in 2019 and 2020.

Windows Systems - The Good, The Bad and The Ugly

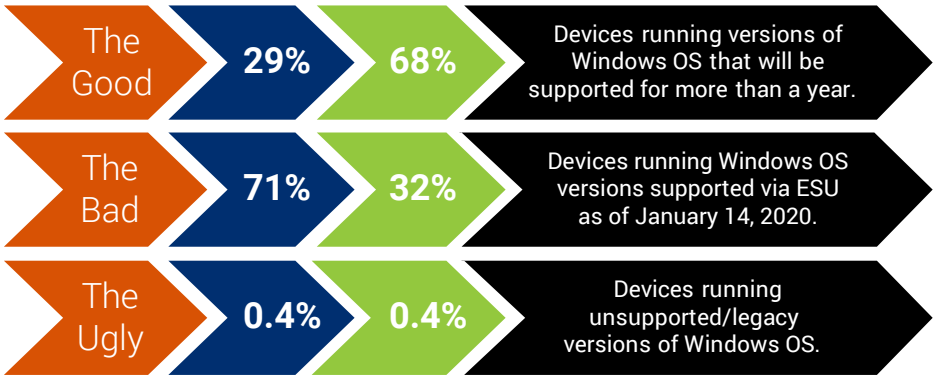

Figure 3 – Legacy Windows instances in 2019 and 2020.

We saw a **reduction of the percentage of "bad" devices from 71% to 32%**, which is good news and indicates that the end of support had a positive effect on upgrading some systems in healthcare.

However, we also noticed that the **percentage of devices running entirely unsupported versions has not changed, remaining constant at 0.4%**. This includes now-obsolete Windows OSes like Windows XP and Windows Server 2003, which were firmly installed in legacy devices that were intended to be used for decades. Many of these devices were manufactured before device vendors fully grasped the cybersecurity and privacy challenges related to maintaining embedded systems. Although a small percentage of devices, they tend to be some of the most critical devices in an HDO and **the fact that the percentage hasn't changed indicates that the legacy OS problem is expected to continue well into the future**.

> What we said in last year's report remains true: "Networks will most likely continue to have medical devices running legacy operating systems since updates are costly. The downtime associated with an operating system update might not be acceptable for critical-care systems. In addition, certain legacy applications simply will not work on more recent versions of Windows due to lack of support, compatibility, or license schema issues. The business need to run legacy operating systems on medical devices isn't going away any time soon, so these devices **must be segmented appropriately to protect access to critical information and services**."

2.1.2. Number of network segments

Network segmentation is a fundamental measure to **limit the attack surface** in healthcare networks. Segmentation is often **achieved by a combination of techniques** at network Layers 2 and 3, including VLANs, Access Control Lists (ACLs), subnetting and firewalling.

To understand at a high level the **use of segmentation in healthcare networks**, we decided to investigate the percentage of deployments running VLANs with medical devices. Again, we compare the data from 2019 with the data from 2020.

Figure 4 shows the **percentage of deployments running between 1 and 150+ different VLANs with medical devices**. We observe a **sharp decrease in deployments running only one VLAN** while there is **some increase in deployments with more than 25 VLANs**.

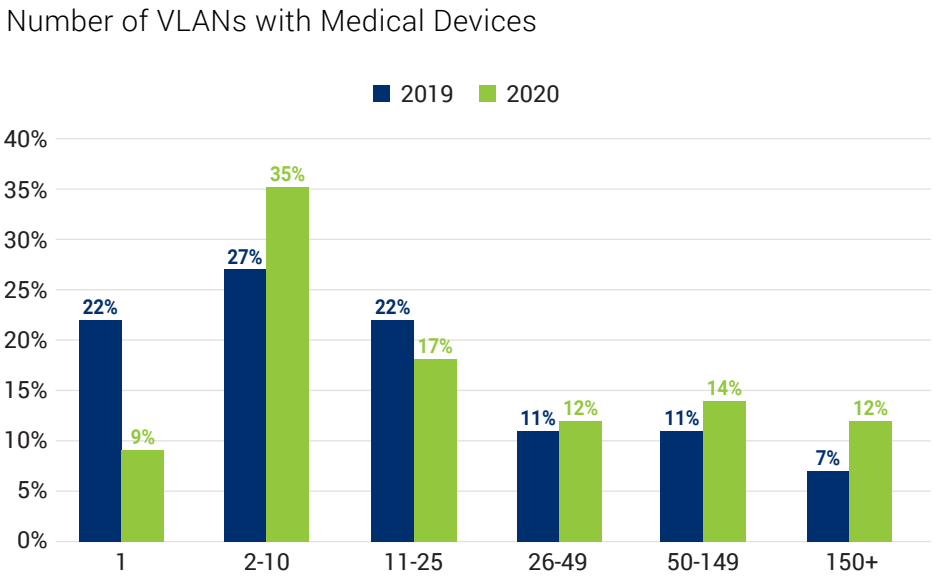

Figure 4 – Number of VLANs with medical devices in 2019 and 2020.

These numbers indicate a **trend toward increasing use of segmentation in healthcare networks**, where segments contain fewer devices and (hopefully) devices that are grouped according to their purpose. To test this last hypothesis, this year we decided to **analyze in more detail the types of devices we see in VLANs in our Device Cloud**.

2.2. Device cloud findings: persistent network segmentation issues

To illustrate how well HDOs are segmenting their networks and controlling configurations, we sought to answer two device and network configuration questions:

1. What is the **mix of IT, IoMT, IoT and OT** devices in production VLANs? (Section 2.2.1)
2. To what extent are devices in these VLANs properly configured? (Section 2.2.2)

2.2.1. Devices on mixed-use VLANs

Even if there are many VLANs in a network, there may be **segments that mix sensitive and vulnerable devices, which means that a vulnerable device may be used to reach a sensitive one**. Below, we focus on observing what types of devices are present in the various healthcare segments and how the mix of different device types can lead to network vulnerabilities.

- For every VLAN with at least one healthcare device, 60% of HDOs also had non-healthcare devices on the same segment. Ninety percent of VLANs have a mix of healthcare devices and IT devices. These numbers indicate that even though the number of

VLANs is increasing, many still do not take device purpose into consideration when designing network segments, which is clearly problematic.

- One of the most detrimental mixed-use examples is that **computers and printers are often present in the same VLAN as healthcare equipment** (e.g., patient monitors, X-ray machines, etc.). Computers in the pharmacy or doctors' workstations may also figure in this mix. **Networking devices like serial-to-IP converters used to connect serial healthcare devices to computing workstations are also on the same segment as these devices**. It is important to note that the security status of the general-purpose computing equipment can directly affect the security status of the specialized healthcare devices communicating on the same VLAN.
- What is more concerning is that we see instances of **personal devices (such as smartphones, smartwatches, tablets) and OT devices on the same VLAN as sensitive healthcare equipment**. These devices might contain vulnerable software or targeted malware which can make other devices on the VLAN susceptible to infection as well.

Below, we provide two observed **examples of poor segmentation in healthcare networks** to illustrate the kinds of device mixes that we mention in the points above.

Faulty segmentation example 1:

Figure 5 shows a VLAN containing a **smartphone, a tablet, a barcode scanner and an infusion pump**. We can immediately see why this VLAN design is problematic. Ideally, the mobile devices should be moved to a network specifically for personal or guest devices, whereas the barcode scanner should be moved to a VLAN dedicated to facilities or equipment scheduling. These Layer 2 segments should then be restricted in their network access privileges to the sensitive medical device segment.

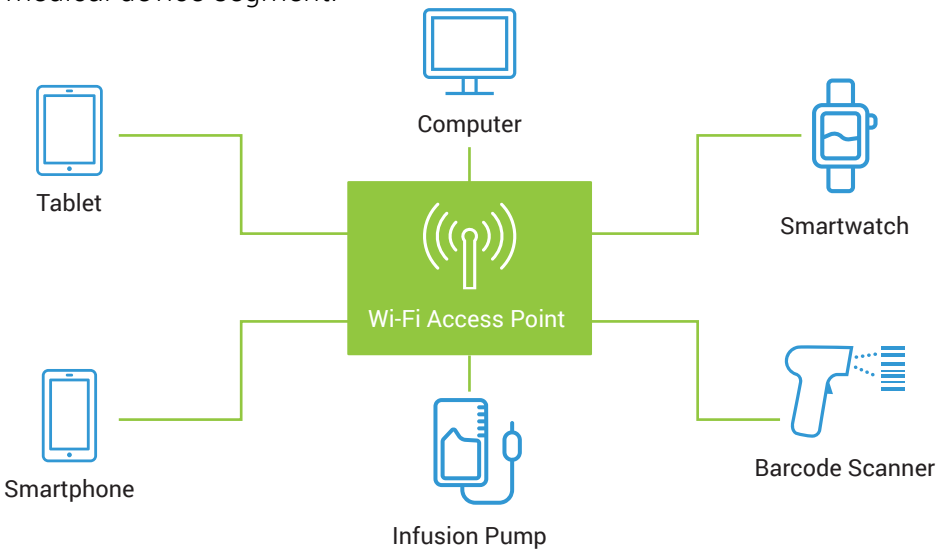

Figure 5 – Example of VLAN shared by smartphones, tablets, barcode scanners and healthcare devices.

Faulty segmentation example 2:

In another healthcare network, as seen in **Figure 6**, we observed a VLAN that has **a printer, a computer, an ultrasound device, a serial-to-IP converter and a wireless clock**. The computer and the printer can be associated with the ultrasound machine, and the serial-to-IP converter may be used to connect a separate healthcare device. However, the clock is unrelated to the 'purpose' served by the VLAN. In this case, the clock may be improperly patched (as are many IoT devices) and hence vulnerable. Therefore, it represents a weakness in the network segment that can potentially affect the other critical devices in the segment. To neutralize the threat to the ultrasound machine and other sensitive associated devices in the VLAN, the clock should be moved to a separate VLAN and firewalled off.

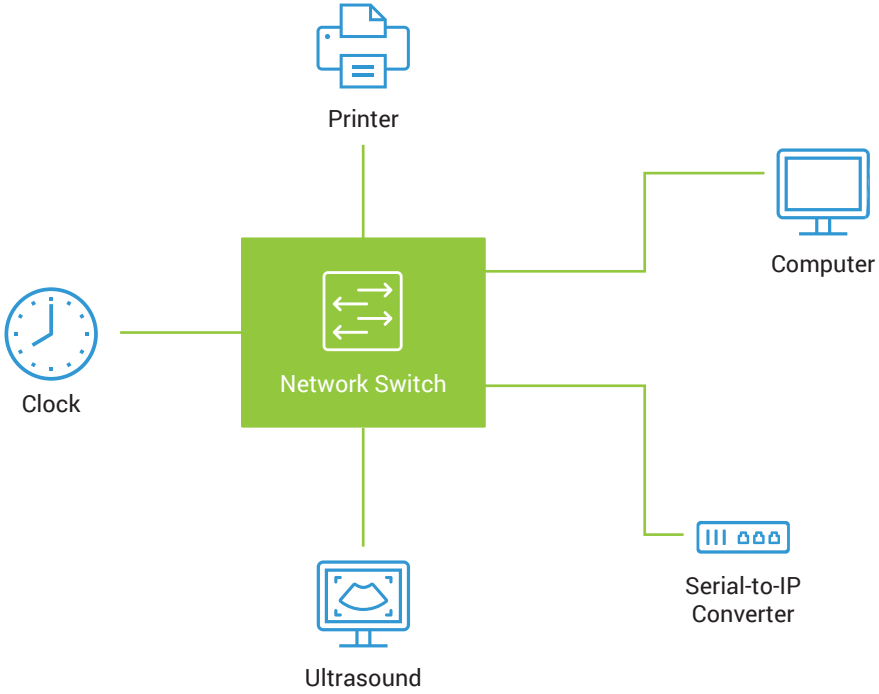

Figure 6 – Example of VLAN shared by a wireless clock and an ultrasound.

2.2.2. Devices with default passwords

Even if a VLAN is properly designed based on device purpose and sensitivity, **having poorly configured devices with default passwords can compromise the segment**.

The OWASP IoT project lists "Weak, Guessable or Hardcoded Passwords" as the top cyber risk to IoT since 2018 [11]. There are a surprising number of IoT devices with default passwords on production networks. If compromised, these devices may act as the springboard for the attacker and can allow immediate lateral movement within the VLAN. Therefore, it is important to proactively identify such vulnerable IoT equipment in a network.

We also identified **healthcare equipment (specifically patient monitors and CT scanners) with default credentials alongside other IT and IoT equipment**. In these scenarios, the healthcare devices act as the weak links on the network. Such an example is observed in **Figure 7**, where a CT scanner with default credentials causes the entire network to be vulnerable. Thus, it is not enough to properly segment the network. Proactive awareness about the device's current security status and its configuration is paramount. Because of that, **we need to dynamically assign devices to segments based on their security status and purpose on the network**.

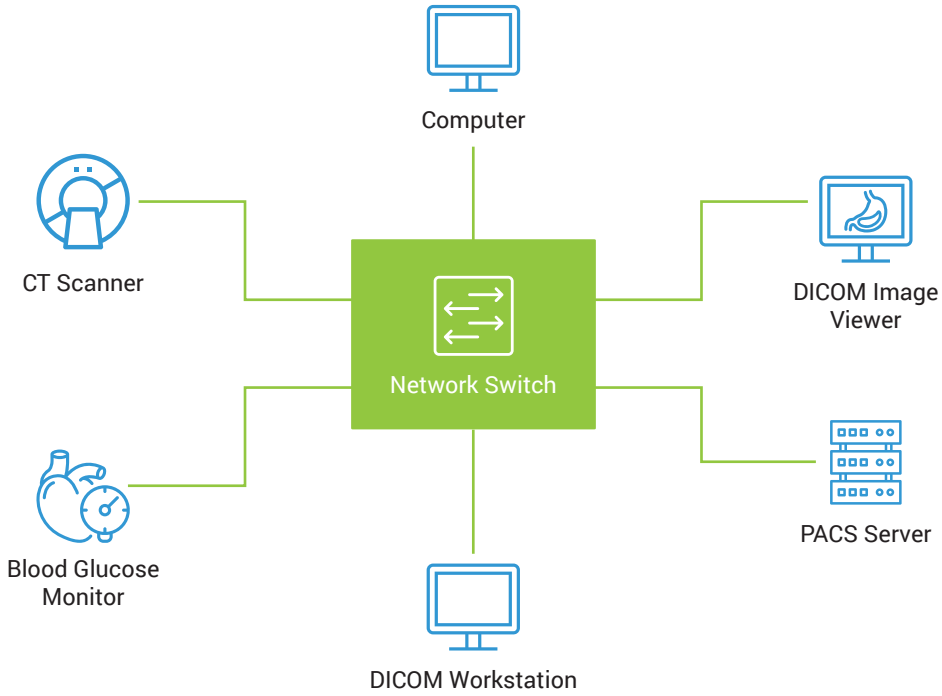

Figure 7 – Example of ideal VLAN setup with a problematic CT scanner using a default password.

2.3. Examining insecure communications through network traffic

Given the trends observed in Sections 2.1 and 2.2, we chose to **deepen our analysis by looking into** another source of data that provides **more detailed information** about healthcare networks: **network traffic analysis from production healthcare networks**. We worked with several large HDOs with the goal of **identifying security issues experienced in practice** by security teams in HDOs.

We did a deep dive into five HDOs with diverse medical devices and a wide variety of standard and proprietary medical protocols to get a rich sample set. Below, we report on the high-risk issues we found, which include **external communications** (Section 2.3.1) and use of **insecure protocols**, both medical and non-medical (Section 2.3.2).

2.3.1. External communications

Our analysis identified several high-risk issues:

- **Four out of the five HDOs were communicating between public and private IP addresses using a medical protocol,** HL7, used **to exchange medical information in clear text**, which can easily be read and can leak sensitive patient information such as names, addresses, family information, allergies and test results. We provide more details about the HL7 protocol in the next section.
- **In two out of the five HDOs, medical devices communicated over IT protocols with external servers** reachable from outside the HDO's perimeter. For example, a medical information system was seen communicating externally over Secure Shell (SSH), another reaching a web server over HTTP, and yet another downloading files from an external file server via FTP. In this last instance, the external server was shown on Shodan with more than 25 vulnerabilities. Attackers would have an easy entry point into the network simply by compromising these external servers to serve malicious files (such as remote access tools).
- We even found that **one out of the five HDOs had an application with electronic health records exposed on the public internet**. This issue has been previously explored by the security community, and it is well known that medical data exposed online is routinely traded by hackers in underground markets [6].

2.3.2. Insecure protocols

Transport Layer Security (TLS) is a cryptographic protocol used to secure network communications of higher-level protocols, such as HTTPS. **Older versions** of this protocol, such as SSLv3, TLSv1.0, and TLSv1.1 are **known to be insecure** and impacted, for instance, by the POODLE and BEAST attacks [18], in which an attacker can downgrade connections and decrypt the traffic, thus being able to access sensitive information. We found that these **insecure versions are still used in all of the HDOs** we analyzed, both internally and externally.

We also found that **all HDOs we analyzed used obsolete versions of other protocols**, such as SNMP versions 1 and 2, used to manage and monitor the status of networked devices, and NTP versions 1 and 2, used to synchronize the clocks of networked devices.

Even more worrisome, we found instances of Telnet in three out of the five HDOs. The clear-text, unencrypted Telnet protocol was designed in 1969 and specified by the IETF in 1983 [19], and has long since been replaced by SSH [20]—but Telnet is still commonly used by devices in HDO networks today.

The most interesting finding is about insecure medical protocols. Medical devices in HDOs transmit data on the network using either **standard medical protocols or proprietary ones**, which are developed by vendors for use within their device ecosystems. Below is a list of some of the most common medical protocols we identified in the HDOs:

- **HL7** [21] is the most widely used **interoperability and data exchange** protocol in medical networks. This messaging standard allows the exchange of patient, clinical and administrative information.
- **DICOM** [22] defines both the format for storing **medical images** and the communication protocols used to exchange them. As a de facto standard, it is implemented by all major vendors of devices involved in medical imaging processes, such as diagnostic workstations, storage servers and medical printers.
- **POCT01** [23] and **LIS02** [24] are used for **point-of-care testing and laboratory testing devices**, respectively. These protocols can issue test orders with patient information to devices and are used by the devices to communicate the results of tests back to a data management system. LIS02 is a revision of a previous standard called ASTM E1394, which was mostly used for serial communication. POCT01, on the other hand, is a newer XML-based protocol.
- We also identified **dozens of proprietary protocols** used in HDOs by equipment such as ventilators, dialysis machines, infusion pumps and patient monitors. These protocols are used by **major vendors** such as Philips, General Electric (GE), Beckton Dickinson (BD), and others.

While supporting critical operations in healthcare delivery organizations, these medical protocols often lack encryption and authentication, or they do not enforce its usage. The HL7, DICOM and POCT01 standards cite the possibility of encrypting transmitted data in their standardization documents but leave the choice of implementation to individual deployments (sometimes assuming that encryption happens at a lower layer, e.g., by using TLS). This is often due to resource constraints in medical devices or the belief that communications in internal "closed" networks do not need to be protected, which is against the modern security mindset of assuming that breaches are inevitable.

The result is as expected: **none of the HDOs analyzed were encrypting HL7, DICOM, POCT01 or LIS02 traffic. The traffic of identified proprietary protocols was also seen in clear text**.

This situation is very similar to what we observe for OT and IoT devices used in Industrial Control Systems or building automation, for instance [13] [25] [26], which **allows attackers to sniff, tamper with and inject malicious traffic into the network**. However, in healthcare networks the potential consequences are much more dire, since the data being transmitted is very sensitive, and the effects of tampering with commands issued by these devices can result in loss of life.

Below are some **examples of sensitive data seen in clear text** in the network traffic via POCT01, LIS02 and a proprietary protocol used by BD Pyxis MedStation medication dispensing systems. **Figure 8** shows LIS02 traffic containing the name and date of birth of two patients and the results of their tests. **Figure 9** shows POCT01 traffic also containing personal data of patients, coming from a Roche Accu-Chek glucose monitor. **Figure 10** shows traffic using the proprietary BD protocol and displaying **patient and doctor information, as well as prescribed medication and quantities dispensed** for each patient. Because of the sensitive nature of the data, personal information has been partially redacted.

Figure 8 – Patient data in clear text (LIS02 protocol).

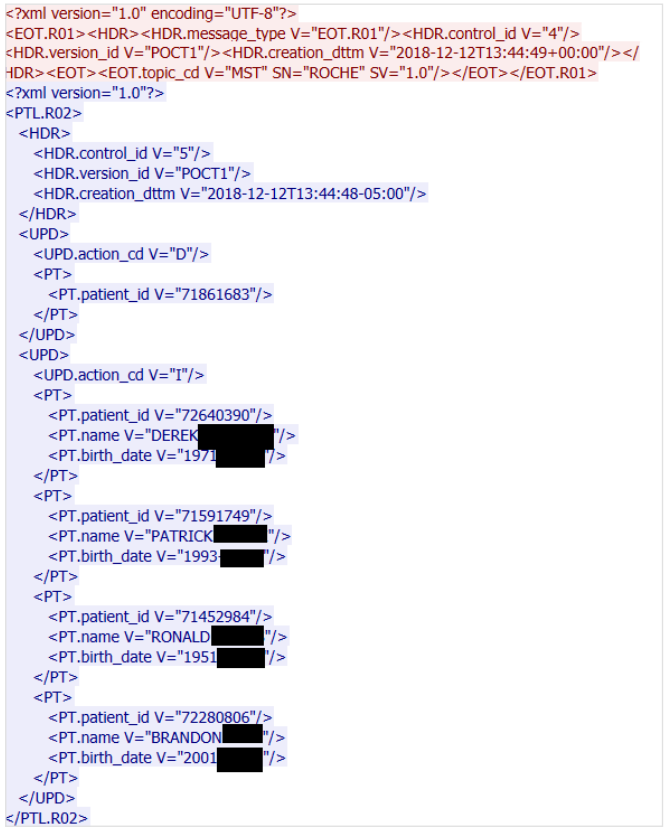

Figure 9 – Patient data in clear text (POCT01 protocol).

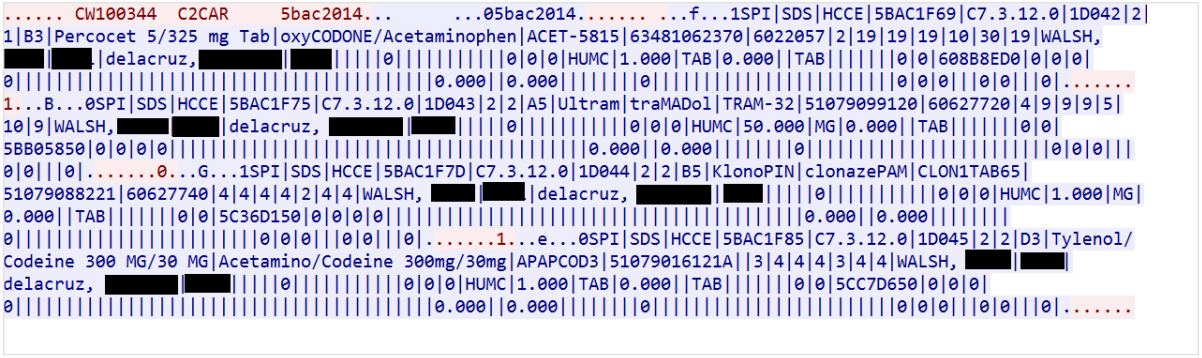

Figure 10 – Patient data in clear-text (BD protocol).

Besides allowing the sniffing of sensitive data, these protocols allow for other attacks, which we will explore in Section 3.

3. Attacking healthcare networks

Cyberattacks may impact the c**onfidentiality, integrity, or availability** – known as the CIA triad – of HDO data, devices, and networks. Lack of confidentiality is an information security and privacy issue, which can have adverse effects on **compliance** efforts. Integrity and availability deficiencies, on the other hand, can have adverse **safety** effects, which in the case of connected medical devices, translates into the possibility of harming patients.

In this section, we first discuss a few **known attacks on HDO networks** (Section 3.1), using some examples reported on the news and some described in the research literature; then, we **demonstrate new attacks that we developed in our lab** (Section 3.2) to illustrate issues with the use of insecure protocols and devices in poorly segmented networks.

3.1. Known attacks

A growing number of data breaches affect the confidentiality of data about patients and employees of HDOs [27] [28] [29]. Besides smaller isolated incidents, in recent years, there have been well-known Advanced Persistent Threat (APT) groups systematically targeting the healthcare industry. For instance, the Orangeworm group has targeted HDOs across Asia, Europe and the United States since 2015 [30]. By leveraging the Kwampirs malware, they were able to infect medical devices such as X-ray and MRI machines. Another group, dubbed APT41, has been stealing data from medical device and pharmaceutical companies since 2014 [31].

Ransomware, which increasingly targets healthcare organizations [32], **severely affects the availability of devices and data, thus potentially paralyzing HDOs** [33]. Ransomware typically affects **unpatched IT devices**, such as workstations, but it may also affect **medical devices** running off-the-shelf operating systems [34]. One of the most famous examples of ransomware was the worldwide WannaCry attack in 2017, which cost the UK's National Health Service alone approximately $100 million [35]. A more recent example is the attack on Fresenius, Europe's largest private hospital operator, which happened in the middle of the COVID-19 pandemic [36].

On the other hand, **attacks affecting the integrity** of data and devices in healthcare have been mostly left out of real incidents, with few mentions in the research literature [37]. However, these attacks are alarming because they may **affect the health of patients by tampering with critical information** (such as allergies, pre-existing conditions and vital sign readings, etc.) **or by disrupting the normal behavior of a medical device** (such as making an infusion pump deliver too much or too little fluid into a patient's body [38] [39]). Attacks of the latter variety usually **target single vulnerable devices**, and there are growing numbers of vulnerabilities disclosed for medical devices [40]. Attacks of the first variety, though, can **leverage the existing insecure communications in HDO networks**.

As discussed in Section 2.3.2, due to the insecurity of many medical protocols, an attacker with access to the network can easily sniff traffic, tamper with data and inject arbitrary packets, thus compromising the integrity of data in healthcare networks. There are some **known attacks** leveraging medical protocols already described in the literature, targeting:

- **HL7**, to tamper with patient medical records and test results, thus allowing attackers to change critical information such as allergies, pre-existing conditions, medication prescriptions and test results [41] [42] [43] [44].
- **DICOM**, to intercept image transfers and change them so that healthy patients show problems, such as tumors, and unhealthy patients show clean scans [45], or to embed executable malware in valid DICOM images [46].
- **Proprietary** protocols, to modify a patient's vital signs (e.g., pulse rate) sent by a GE patient monitor over their RWHAT protocol [47].

These attacks can have multiple outcomes, such as getting someone treatment they do not need or depriving someone of the necessary treatment. These attacks can also be chained and extended with more complex options. **Targeted attacks** can combine some of the options above to target one specific patient. For instance, an attacker could change a patient's EHR allergy entry and maintain his patient monitor reading at a normal level, even if he is given a dangerous medication. Untargeted attacks can replicate the effects of any of the attacks above to multiple patients in a hospital. For instance, an attacker can randomly change the test results of several patients or make several patient monitors show patients flatlining at the same time.

3.2. Reproducing attacks in the lab

To demonstrate in practice the exploitation of communications in a healthcare network, we set up a small healthcare lab (depicted in Figure 11) containing:

- On the clinical side, a Philips IntelliVue MP50 patient monitor and a Siemens DCA Vantage blood and urine analyzer (a common point of care testing device).
- On the IT side, a Central Monitoring Station (CMS) that shows the real-time readings of the patient monitor using IxTrend Express software and a Laboratory Information System (LIS) that stores test results from the blood analyzer. Central Monitoring Stations are common in hospitals to remotely display the result of several patient monitors. Since we did not find a suitable and simple open-source solution for LIS, we implemented a simple LIS02 server using Python ASTM and a POCT01 server according to the communication specifications of the Siemens DCA Vantage analyzer [48].

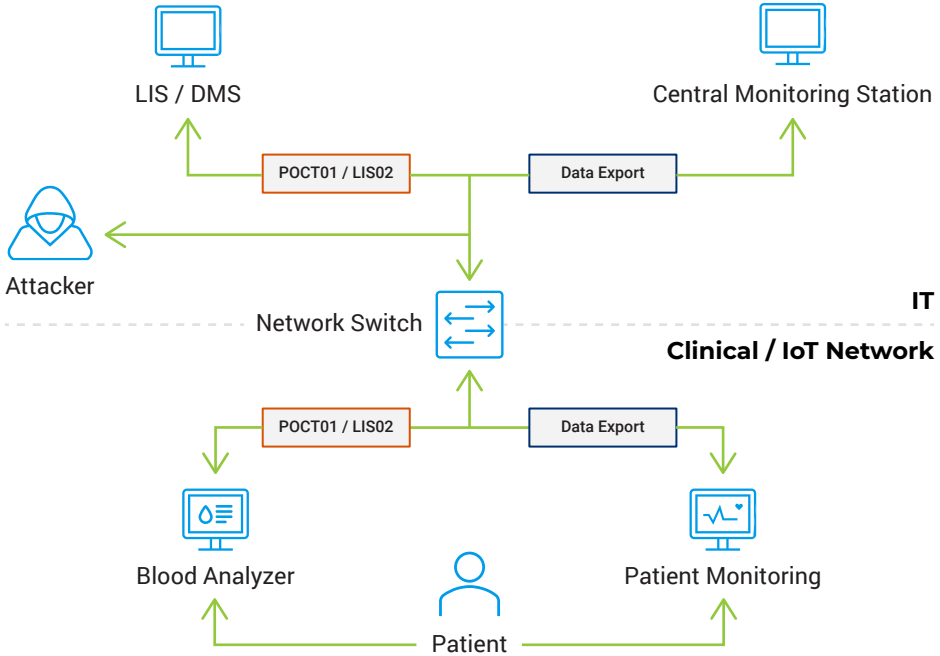

Figure 11 – Healthcare lab used for our attack demonstrations.

All devices in this lab are connected to the same network switch in the center, and there is no segmentation on the network. Although that is an oversimplification, we discussed at length in Section 2.2 how in practice networks in HDOs are improperly segmented. The result is that sensitive devices are reachable not only from systems used by medical staff but also by guests and adversaries that have physical access to the network.

The ease of obtaining physical access is common in hospitals since network sockets are often used in patient rooms to connect medical devices that interface directly with patients [14] [45]. Therefore, we also represent in Figure 11, an attacker who has local access to the network. Remote access to an HDO's network, on the other hand, is most commonly achieved by attackers via phishing [9].

The attacks that we chose to demonstrate in the following sections are similar to the ones leveraging well-known protocols described at the end of Section 3.1. We use different protocols than the ones already described in the literature (HL7, DICOM, and RWHAT) to show that those issues are prevalent in healthcare networks. We implemented our attacks using the protocols POCT01, LIS02 and Philips Data Export, as described in **Table 1**. While POCT01 and LIS02 were briefly described in Section 2.3.2, Data Export is a protocol used by Philips patient monitors, such as the IntelliVue MP50. This protocol allows patient monitors to communicate vital readings to a central monitoring system and otherwise aggregate the readings of multiple patients. It is somewhat similar to the GE RWHAT protocol mentioned in Section 3.1.

Table 1 - Attacks implemented in the lab.

Attack Example #	Type of attack	Protocol	Description
1	Confidentiality	POCT01	Dump test results stored in a device via POCT01
2	Integrity	LIS02	Change test results sent from a device to the LIS via LIS02
3	Availability	Data Export	Abort connection between patient monitor and CMS
4	Integrity	Data Export	Change real-time pulse reading shown in the CMS

Before we discuss the attacks we implemented below, some quick observations are in order:

1. Both medical devices used in our experiments (Philips IntelliVue MP50 and Siemens DCA Vantage) were obtained via an online auction site and were pre-owned.

2. Although we did not perform a thorough vulnerability assessment of the devices, we spotted two immediate issues on the DCA Vantage: the possibility of breaking out of kiosk mode and running code on the device (which runs a full-fledged Windows CE); and the use of a hard-coded password for the database containing sensitive data such as administrator passwords, test results and device logs. These issues are not used in the attacks below, but were reported to Siemens. These issues were assigned CVE-2020-15797 and CVE-2020-7590, respectively.

3. Issues 1 and 2 above allowed us to retrieve test results from the Siemens DCA that were stored on the device, presumably from a previous owner. Although we did not investigate further the prevalence of sensitive data in second-hand medical devices, this shows that HDOs should pay special attention to data confidentiality when disposing of medical devices.

4. The attack examples mentioned in Table 1 represent the goals of an attacker who wants to disrupt a healthcare network. This is the focus of what we want to present in the next sections, so in our lab setup, we adopted an attacker model that assumes an attacker is already inside the network with the ability to sniff and, when necessary for the attack, modify packets in the network, essentially acting as a man-in-the-middle (MitM). The most popular way of achieving MitM in a network is via ARP poisoning

(also known as ARP spoofing). The Address Resolution Protocol (ARP) is used by devices in a network to resolve IP addresses to physical MAC addresses. ARP poisoning exploits the lack of authentication in the protocol by sending spoofed messages to the network with the goal of associating the attacker's host MAC address with the IP address of a target host. This can be achieved automatically using tools such as Ettercap software.

3.2.1. Attack example 1: Dumping test results

The goal of this attack is to intercept test results being sent from the DCA Vantage analyzer to the LIS, although this attack (at least the passive variant) could be reproduced with any two devices communicating over unencrypted and unauthenticated POCT01.

An example of a blood test result from the DCA Vantage device is shown in **Figure 12**. It shows the result of a Hemoglobin A1C test. The result shown in Figure 12 is 66 mmol/mol, which could be indicative of diabetes.

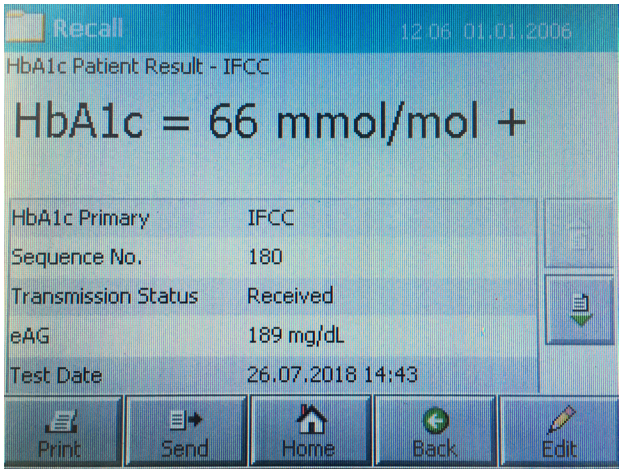

Figure 12 – HbA1c test result shown on the screen of the DCA Vantage.

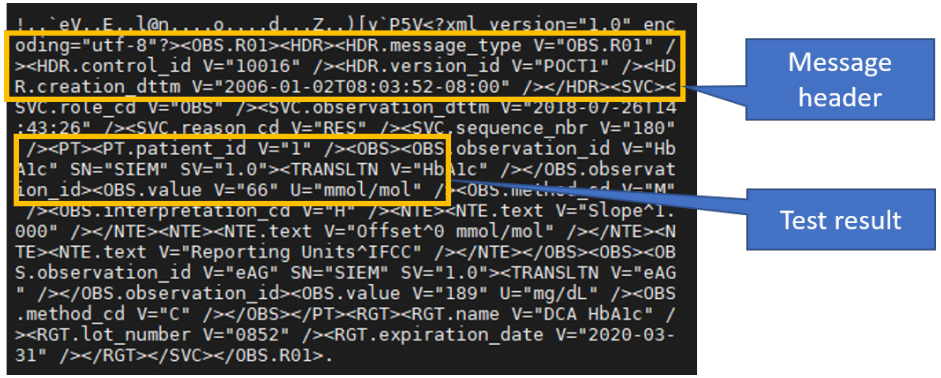

Figure 13 – Details of the POCT01 packet transmitting the HbA1c test results to the LIS.

When the operator chooses to send a test result (Figure 14) to the LIS via the POCT01 protocol, an established synchronous POCT01 conversation exists between the DCA Vantage and LIS, a packet such as the one shown in **Figure 14** is generated and sent over the network. The packet contains a header that specifies the message type ("OBS.R01" stands for a test result sent from DCA), the timestamp, the header, as well as the test result creation timestamp (the result we use here was stored in the device, and is from 2006). Further, the packet lists the patient ID and the test result value.

Since the above packet is transmitted in clear text, attackers can **passively intercept test results** sent over by operators by simply sniffing the network traffic and examining the POCT01 packets that contain the "OBS.R1" message type in the message header.

However, attackers can also **actively intercept test results** by bringing rogue devices that can serve as **fake LIS servers** into hospitals. Due to the lack of traffic encryption, these devices can then hijack communications between a POCT device and a legitimate LIS server (e.g., via ARP cache poisoning). Then, attackers can execute a limited set of remote commands via POCT01 that the

device supports: e.g., force the device to send all pending test results to the fake LIS server, or update the list of device administrators. Considering the state of the network segmentation in medical networks that we observed, this is a realistic scenario.

As a proof-of-concept, we have implemented a fake LIS server according to the device-specific POCT01 communication protocol implemented in the DCA Vantage [48]. We first perform an ARP cache poisoning attack with Ettercap so that the DCA Vantage is forced to communicate with our proof-of-concept server. Once the device sends the hello message ("HEL.R01"), our server responds with an ack message ("ACK.R01"), requests pending tests results ("REQ.R01") and obtains them, and, after a short conversation sequence (detailed in [48]) establishes a continuous conversation mode with the DCA Vantage. In this mode, all further test results will be sent directly to the fake LIS server. Moreover, the DCA Vantage will accept a limited set of commands from the server, such as to update the list of the device's operators ("OPL.R01").

3.2.2. Attack example 2: Changing test results

The goal of this attack is to tamper with a test result being sent from the DCA Vantage analyzer to the LIS via the LIS02 protocol, although this attack could be reproduced with any two devices communicating over unencrypted and unauthenticated LIS02 or POCT01.

When the operator chooses to send a test result (the same one seen in Figure 12) to the LIS via the LIS02 protocol, a packet such as the one shown in **Figure 14** is generated and sent over the network.

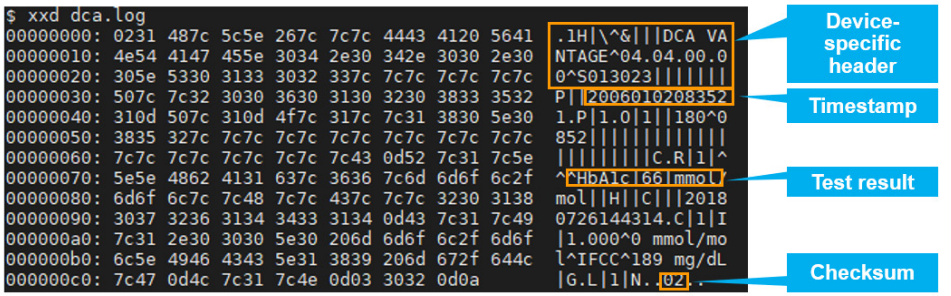

Figure 14 – Details of the LIS02 packet transmitting the HbA1c test result to the LIS.

Notice that the packet contains a header with some information about the device issuing the result (DCA Vantage), a timestamp, detailed test results and a checksum at the end. (For a complete reference on the contents of a LIS02 packet, see [24]).

When the LIS server receives the packet, it displays the results, as shown in **Figure 15** and stores them internally.

```
----------------------------
RECEIVED!
----------------------------
HEADER: ['H', [[None], [None, '&']], None, None, ['DCA VANTAGE', '04.04.00.00', 'S0113023'],
 None, None, None, None, None, None, 'P', None, '2006010208352']

PATIENT: ['P', '1']

ORDER: ['O', '1', None, ['180', '0852']], None, None, None, None, None, None, None, None, No
ne, None, None, None, None, None, None, None, None, None, None, None, None, 'C']

RESULT: ['R', '1', [None, None, None, 'HbA1c'], '66', 'mmol/mol', None, 'H', None, 'C', None
, None, '20180726144314']

COMMENT: ['C', '1', 'I', ['1.000', '0 mmol/mol', 'IFCC', '189 mg/dL', 'G']]

TERMINATOR: ['L', '1', 'N']
```

Figure 15 – HbA1c test result received by the LIS.

When an attacker wants to tamper with this flow and send incorrect results to the LIS, he must do the following: create a man-in-the-middle, drop the original packet, create a new packet with a modified test result, compute the new checksum, insert it into the new packet and send it. This can be easily achieved with a tool such as Ettercap and a custom filter that modifies test results and checksums.

The result of a modified packet received by the LIS server is shown in Figure 16. This would cause the LIS to store an incorrect result. In this example, the attacker has changed the results of a diabetic patient to a normal test result (41 mmol/mol). The opposite could just as easily be done.

```
----------------------------
RECEIVED!
----------------------------
HEADER: ['H', [[None], [None, '&']], None, None, ['DCA VANTAGE', '04.04.00.00', 'S0113023'],
 None, None, None, None, None, None, 'P', None, '2006010208357']

PATIENT: ['P', '1']

ORDER: ['O', '1', None, ['180', '0852']], None, None, None, None, None, None, None, None, No
ne, None, None, None, None, None, None, None, None, None, None, None, None, 'C']

RESULT: ['R', '1', [None, None, None, 'HbA1c'], '41', 'mmol/mol', None, 'H', None, 'C', None
, None, '20180726144314']

COMMENT: ['C', '1', 'I', ['1.000', '0 mmol/mol', 'IFCC', '189 mg/dL', 'G']]

TERMINATOR: ['L', '1', 'N']
```

Figure 16 – LIS displaying changed HbA1c test result due to an attack.

3.2.3. Attack example 3: Disconnecting a patient monitor

The goal of this attack is to close an ongoing connection between the patient monitor and the CMS so that the staff remotely monitoring a patient loses the real-time information about vital readings.

This is achieved by issuing an "Association Abort" command that is available on the Data Export protocol, as shown in **Figure 17** (which comes from page 337 of the protocol manual [49]).

ASSOCIATION ABORT

The following building blocks can be used to format a Association Abort message:

AbortSessionHeader

```
0x19 0x2E
```

AbortSessionData

```
0x11 0x01 0x03
```

AbortPresentationHeader

```
0xC1 0x29 0xA0 0x80 0xA0 0x80 0x30 0x80
0x02 0x01 0x01 0x06 0x02 0x51 0x01 0x00
0x00 0x00 0x00 0x61 0x80 0x30 0x80 0x02
0x01 0x01 0xA0 0x80 0x64 0x80 0x80 0x01
0x01 0x00 0x00 0x00 0x00 0x00 0x00
```

AbortUserData

This block is empty in the Abort message.

AbortPresentationTrailer

```
0x00 0x00 0x00 0x00
```

Figure 17 – Association Abort message described in the Philips Data Export manual [46].

The attacker can spoof an abort message and make the CMS believe that the monitor wants to close the ongoing connection, thus causing a denial of service. To do so, the attacker simply has to craft a packet containing a payload with the abort message and send it to UDP port 24105, which is used by default the Data Export protocol (or any other port that is used in the connection between the monitor and the CMS, which can be learned by sniffing the traffic). The contents of the payload are as shown in **Figure 17**: 0x19 0x2e 0x11 0x01 0x03 0xc1 0x29 0xa0 0x80 0xa0 0x80 0x30 0x80 0x02 0x01 0x01 0x06 0x02 0x51 0x01 0x00 0x00 0x00 0x00 0x61 0x80 0x30 0x80 0x02 0x01 0x01 0xa0 0x80 0x64 0x80 0x80 0x01 0x01 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00.

The result of the attack can be seen in **Figure 18**, where the CMS is shown displaying an error message informing the user that the monitor has closed the connection and stopped sending data to the CMS.

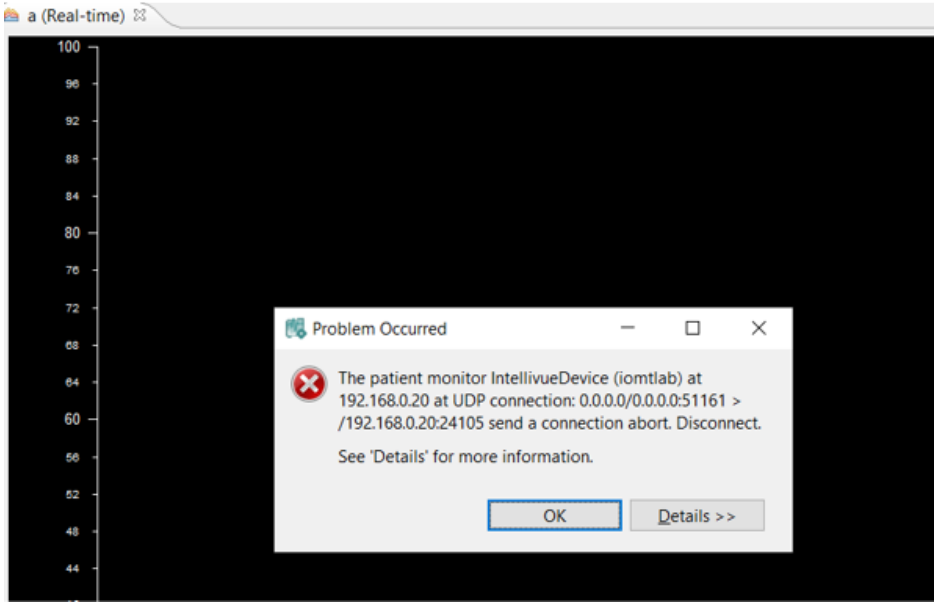

Figure 18 – CMS displaying the result of an Association Abort message.

3.2.4. Attack Example 4: Changing a patient's vital readings

The goal of this attack is to tamper with the vital readings sent from the patient monitor to the CMS so that the staff remotely monitoring a patient sees incorrect real-time information about vital readings.

This is achieved by modifying on-the-fly the Data Export packets sent from the monitor to the CMS. To do so, the attacker can again use Ettercap and create a filter that replaces the real-time vital readings with a desired value.

The only challenge, in this case, is to understand at which offset in the packets, the vital readings are encoded, since, as shown in the previous section, Data Export is a binary protocol. This information can be obtained from the Data Export manual on page 118 [49], as shown in **Figure 22** for the pulse rate (which we use in the examples below). The manual lists other vital signs, such as blood pressure and oxygen saturation.

Pulse	Pulse Rate from Plethysmogram	
	Label:	
	NLS_NOM_PULS_OXIM_PULS_RATE	0x00024822
	Observed Value:	
	NOM_PLETH_PULS_RATE	0x4822
	Units:	
	NOM_DIM_BEAT_PER_MIN	0x0AA0

Figure 19 – Patient's pulse rate encoding shown in the Philips Data Export manual [46].

Searching for the values 0x4822 and 0x0aa0 on captured traffic between the monitor and the CMS, we find what is shown in the UDP packet in **Figure 21**, where the bytes with values 48 22 indicate to the CMS that a pulse value is incoming and the bytes with values 0a a0 indicate the unit (beats per minute). Finally, the last two bytes encode the actual value of the pulse observed in the monitor, which in this case is 50 in hexadecimal or 80 in decimal. Therefore, we can calculate the offset of the byte we want to change (the one with value 50).

```
02e0  42 28 00 00 00 00 09 24  00 04 00 02 48 22 09
02f0  00 10 00 0e 00 50 00 75  00 6c 00 73 00 65 00
0300  00 00 09 11 00 02 00 06  09 50 00 0a 48 22 00
0310  0a a0 00 00 00 50 83 a7  00 07 00 4e 09 21 00
0320  83 a7 09 2f 00 04 00 01  00 06 09 3f 00 0c 00
0330  20 00 00 01 42 28 00 00  00 00 09 24 00 04 00
```

Figure 20 – Packet capture showing the patient's pulse rate transmitted from the patient monitor.

Once this offset is discovered, the attacker can create an Ettercap filter to extract the right packet containing the patient data and modify the pulse value of the patient to his desired value (e.g., 0 to simulate a patient flatlining or a rapid succession of high and low numbers to simulate an arrhythmia condition) and forward it to the CMS to display this information.

Figure 21 shows the actual reading on the patient monitor, which is a normal pulse of 83. **Figure 22** shows the result of the flatlining attack, as seen by staff on the CMS. Notice that the pulse suddenly drops from a normal range between 70 and 80 to 0. **Figure 23** shows the result of the arrhythmia attack, again as seen on the CMS. Notice that the pulse suddenly increases from the normal range to an accelerated value of 100, and after a short period drops to a low value of around 50 and then repeats the pattern many times. Similar attacks could be implemented to change oxygen saturation, blood pressure and other readings.

Figure 21 – Normal pulse rate reading on the patient monitor.

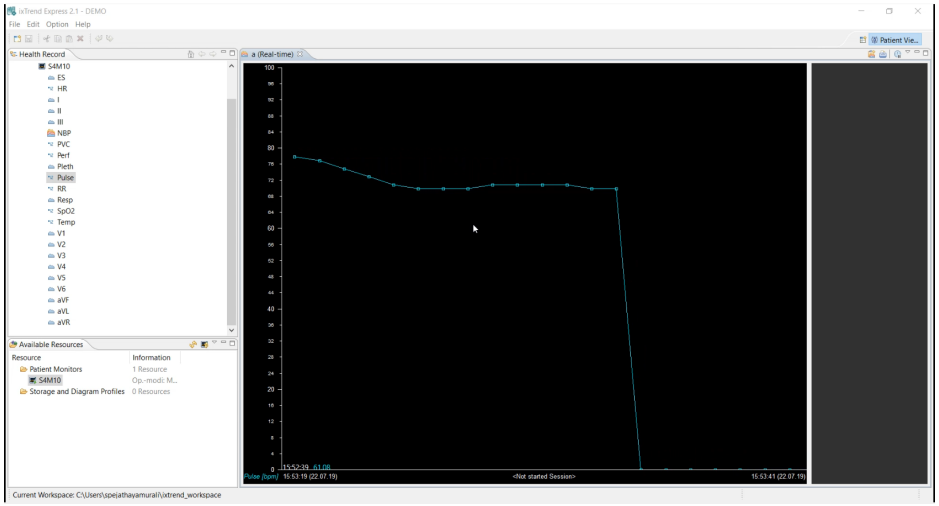

Figure 22 – Result of a flatlining attack shown in the CMS.

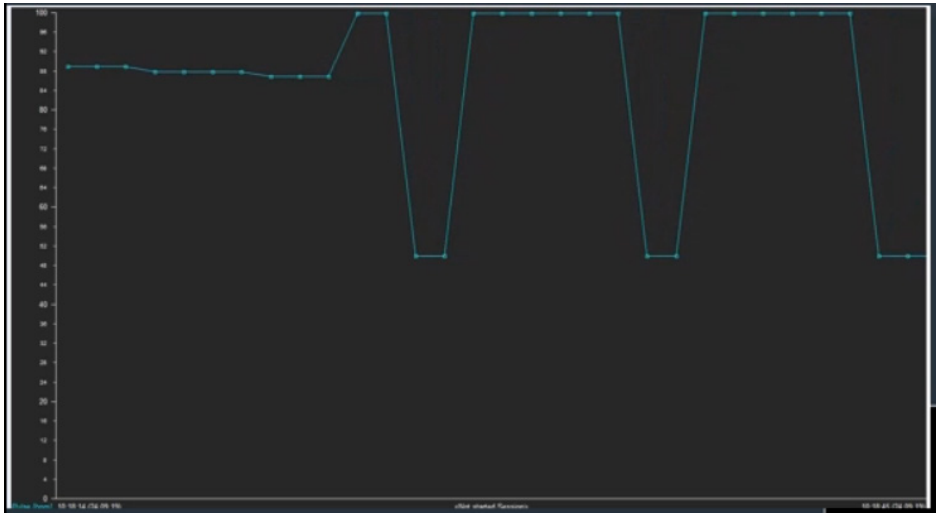

Figure 23 – Result of an arrhythmia attack shown in the CMS.

4. Defending healthcare networks

After analyzing and describing how to attack healthcare networks in the previous sections, we discuss **effective strategies to defend healthcare networks** from cyberattacks. To present these strategies, we make use of the Zero Trust Architecture framework as a guide and describe how to achieve device visibility (Section 4.1), network segmentation and policy enforcement (Section 4.2), as well as automation and orchestration (Section 4.3).

4.1. Get complete visibility into all connected devices and their risk

It is widely recognized that cybersecurity for medical devices will be challenging for the next 20 years, and that visibility is the key to improvement [50].

Leading technology vendors and analysts agree that visibility is foundational. That's why visibility is explicitly baked into defensive frameworks like Forrester's Zero Trust security model [51], while the Center for Internet Security places visibility – in the form of Asset Inventory for Hardware and Software – as the first of twenty critical controls [52].

The importance of visibility in defense, put simply by Forrester analysts, is that "you can't combat a threat you can't see or understand" [51]. All leading interpretations of Zero Trust place visibility as foundational to resource defense. Often, device security comes first in practical discussion of technical controls, as in Google's interpretation: "Device and host inventory is the primary prerequisite to any inventory-based access control" [53]. VMware also places Device Trust first [54], while Microsoft posits that Identities and Devices come together initially as the subject of policy enforcement [55].

Forescout agrees with those views and also believes that visibility is foundational for security, especially for IoT devices, which are so hard to manage. That is why the Forescout platform takes a visibility-first approach to deliver granular, contextual insights into customers' entire device landscapes without disrupting critical business processes. After discovering connected devices, the Forescout platform auto-classifies and assesses those devices against company policies. The combination of these three capabilities – discovery, classification and assessment delivers the device visibility to drive appropriate policies and actions. Notice that this visibility can be configured to be completely passive to avoid the risk of disrupting critical devices attached to patients.

But visibility must extend beyond users and devices. In Zero Trust, the next essential layer is network visibility, with a focus on transport and session security. That is where techniques may be implemented to detect anomalous and malicious network behavior and defend against the attacks demonstrated in the previous section.

Detecting anomalous behavior is challenging on healthcare networks due to their diverse, heterogeneous nature. This variety of devices, applications and protocols adds complexity to network monitoring. For instance, protocol-based detective controls are a must for granular transport and session security policy. Protocol-sensitive deep packet inspection (DPI) is the table stakes required for intrusion detection on HDO networks.

The Forescout solution provides in-depth visibility and cyber resilience with not just assets but also communications inventory based on DPI for IT, OT and healthcare protocols. This allows for network monitoring and threat hunting capabilities, such as threat and vulnerability indicators.

4.2. Implement network segmentation to reduce likelihood and impact of breaches

Segmentation is a Zero Trust principle covered in recent guidance by NIST and ENISA. Further, Gartner analysts suggest enterprises that isolate/segment their campus network devices will experience 25% fewer successful cyberattacks.

Segmentation is such a fundamental control that it impacts every component of the Zero Trust Architecture. Segmenting flat networks while allowing patient data flow via central EMR systems is just one of many HDO use cases for granular segmentation security policy. Example segmentation policies may be driven by the need to:

- Classify and control a diverse array of devices by function and vendor (Device Security)
- Enable an extended workforce, remote vendor support and business associates (User Security)
- Isolate fragile legacy applications and operating systems (Workload Security)
- Protect sensitive patient data stores (Data Security)
- Safeguard the availability and reduce exposure to critical applications that save and sustain lives (Network Security)
- Enable reactive security policy enforcement via Automation & Orchestration

These policies can be developed by a top-down approach using business logic or technically in a bottom-up approach. Such bottom-up opportunities include "good traffic" candidates for whitelisting and "bad traffic" candidates for blacklisting. This consideration is useful in practice, especially in HDO environments with many internal departments and an array of internal business services. The table below presents examples of security policy considerations for granular segmentation.

Good Traffic	Bad Traffic
• DICOM Workstation to PACS • MRI & UltraSound to PACS • DICOM Workstation to MRI & UltraSound & PACS • Radiology to PACS • Nursing & Radiology to EHR • Infusion Pump to Controller	• Non-medical users to EHR • MRI (which runs EoL Windows) accessing regular Windows workstations on standard Windows protocols (Windows being Windows)

Simulating security policy prior to enforcement is crucial – it is one of the only ways to ensure that access to data and applications is appropriately limited without causing downtime, malfunction or other breaking changes, all of which can be disruptive to HDO operations and potentially harm patients or prevent the saving of life itself.

The process of responsibly implementing visibility to simulate policy prior to enforcement typically follows this approach:

1. Create appropriate groups based on the use case and "good traffic" vs. "bad traffic,", e.g., device function, user department, network protocol, application status and data sensitivity.
2. Learn how those groups communicate across the organization and apply filters to identify specific communication patterns and protocols.
3. Simulate policies to tighten communications and adhere to segmentation requirements.
4. Refine policy rules when exceptions are identified.
5. Enable automated response to any policy violations (alert or enforce).

This is how HDOs can responsibly design and deploy security policies to enforce device, user, network, application and data segmentation in their healthcare environments.

The Forescout platform accelerates the design, planning and deployment of dynamic network segmentation across the extended enterprise to reduce your attack surface and regulatory risk. It simplifies the process of creating context-aware segmentation policies and allows visualization and simulation of policies prior to enforcement for proactive fine-tuning and validation.

4.3. Embrace solutions that enable Security Automation & Orchestration (SAO)

When defining Zero Trust, Forrester layered Visibility & Analytics as the foundational prerequisite to the second control layer: *Security Automation & Orchestration (SAO)*. After all, we rely on visibility to design advanced segmentation policies that are enforced by Automation & Orchestration.

In HDO environments, that means integrating visibility across the entire security solutions portfolio. Such integration is required not only to make the policy enforcement control layer work, but also to maximize solutions portfolio ROI.

A cross-portfolio, vendor-agnostic, interoperable solutions architecture is crucial to effective SAO. At the end of the day, orchestration is all about solutions integration. Forrester has suggested that adopting Zero Trust Architecture across the solutions portfolio can reduce an organization's risk exposure by 37% while reducing security costs by 31%.

Forescout eyeExtend shares device context between the Forescout platform and other IT and security products to automate policy enforcement across disparate solutions and accelerate system-wide response to mitigate risks.

5. Conclusion

Our analysis reveals that while HDOs have taken some meaningful steps to better secure their connected devices and networks, there are still several cybersecurity gaps and risks that need to be addressed.

HDOs will have to contend with medical devices running legacy operating systems for the foreseeable future. Hence, it is imperative to identify and mitigate this risk.

Segmentation is a foundational control for risk mitigation in networks with a diversity of IT, IoT and OT devices. However, segmentation requires well-defined trust zones based on device identity, risk profiles and compliance requirements for it to be effective in reducing the attack surface and minimizing blast radius. Over-segmentation with poorly defined zones simply increases complexity without tangible security benefits.

Based on our research, we recommend that HDOs prioritize the following best practices to reduce security and operational risk in healthcare networks:

- **Legacy devices and operating systems.** Accurate identification and classification of medical devices running legacy operating systems are paramount for risk mitigation. Devices that cannot be retired or patched should be segmented appropriately to restrict access to critical information and services only.
- **External communications and exposure.** Network flow mapping of existing communications is not just a prerequisite for designing effective segmentation zones, it also provides a baseline understanding of external and internet-facing communication paths. This can help identify unintended external communications and prevent medical data from being exposed publicly.
- **Insecure and unencrypted protocols.** Start with a network flow mapping project to identify protocols in use. Whenever possible, switch to using encrypted versions of protocols and eliminate the usage of insecure, clear-text protocols such as Telnet. When this is not possible, use segmentation for zoning and risk mitigation.
- **Default, weak or hardcoded passwords.** Identify and remediate weak and default passwords. A single weak link on a network segment can compromise the entire segment. If hardcoded passwords cannot be remediated, leverage segmentation for zoning and isolation.
- **Effective segmentation.** Segmentation can be used as a compensating control and risk mitigation technique for all of the above scenarios. It is also a best practice for compliance ring-fencing, limiting lateral movement and reducing the blast radius of attacks. While there is increasing awareness of the benefits of segmentation, examples of over-segmentation, under-segmentation and poorly designed segmentation zones abound. Start by accurately identifying devices you want to segment by business context and understanding existing network flows between device groups. Then design appropriate zones and access policies to gain the positive security outcomes of segmentation.

REFERENCES

[1] G. O'Brien, S. Edwards, K. Littlefield, N. McNab, S.Wang and K. Zheng, "Securing wireless infusion pumps in healthcare delivery organizations," NIST, 2018. [Online]. Available: https://www.nist.gov/publications/securing-wireless-infusion-pumps-healthcare-delivery-organizations.

[2] J. Cawthra, B. Hodges, J. Kuruvilla, K. Littlefield, B. Niemeyer, C. Peloquin, S. Wang, R. Williams and K. Zheng, "Securing Picture Archiving and Communication System (PACS)," NIST, 2019. [Online]. Available: https://www.nccoe.nist.gov/sites/default/files/library/sp1800/hit-pacs-nist-sp1800-24-draft.pdf.

[3] S. Islam, D. Kwak, M. Kabir, M. Hossain and K. Kwak, "The internet of things for health care: A comprehensive survey," IEEE Access, vol. 3, p. 678–708, 2015.

[4] F. Jaigirdar, C. Rudolph and C. Bain, "Can I trust the data I see?: A physician's concern on medical data in IoT health architectures," in Australasian Computer Science Week Multiconference, 2019.

[5] F. Alsubaei, A. Abuhussein and S. Shiva, "Security and privacy in the internet of medical things: Taxonomy and risk assessment," in IEEE 42nd Conference on Local Computer Networks Workshops, 2017.

[6] M. Fuentes, "Cybercrime and other threats faced by the healthcare industry," TrendMicro, 2017. [Online]. Available: https://www.trendmicro.com/content/dam/trendmicro/global/en/security-intelligence/research/reports/wp-cybercrime-&-other-threats-faced-by-the-healthcare-industry.pdf.

[7] A. Wirth, "The economics of cybersecurity," Biomedical Instrumentation & Technology, vol. 51, no. s6, p. 52–59, 2017.

[8] Carbon Black, "Healthcare cyber heists in 2019," 2019. [Online]. Available: https://www.carbonblack.com/resources/threat-research/healthcare-cyber-heists-in-2019/.

[9] HIMSS, "2019 HIMSS Cybersecurity Survey," 2019. [Online]. Available: https://www.himss.org/2019-himss-cybersecurity-survey.

[10] Forescout, "Putting Healthcare Security Under the Microscope," 2019. [Online]. Available: https://www.forescout.com/wp-content/uploads/2019/05/forescout-healthcare-report.pdf.

[11] OWASP, "Internet of Things Project" [Online]. Available: https://www.owasp.org/index.php/OWASP_Internet_of_Things_Project.

[12] Forescout, "BAS Research Report: The Current State Of Smart Building Cybersecurity," 2019. [Online]. Available: https://www.forescout.com/securing-building-automation-systems-bas/.

[13] Forescout, "Rise of the Machines: Transforming Cybersecurity Strategy for the Age of IoT," 2019. [Online]. Available: https://www.forescout.com/places-in-network/building-automation-system-bas/transforming-cybersecurity-strategy-for-the-iot/.

[14] ISE, "Securing Hospitals: A Research Study and Blueprint," 2016. [Online]. Available: https://www.securityevaluators.com/hospitalhack/.

[15] D. Miessler, "Default passwords," [Online]. Available: https://github.com/danielmiessler/SecLists/blob/master/Passwords/Default-Credentials/default-passwords.csv.

[16] Microsoft, "Windows 7 support ended on January 14, 2020," [Online]. Available: https://support.microsoft.com/en-us/help/4057281/windows-7-support-ended-on-january-14-2020.

[17] Microsoft, "Lifecycle FAQ - Extended Security Updates," [Online]. Available: https://support.microsoft.com/en-us/help/4497181/lifecycle-faq-extended-security-updates.

[18] Y. Sheffer, R. Holz and P. Saint-Andre, "Summarizing Known Attacks on Transport Layer Security (TLS) and Datagram TLS (DTLS)," IETF, 2015. [Online]. Available: https://tools.ietf.org/html/rfc7457.

[19] J. Postel and J. Reynolds, "Telnet Protocol Specification," IETF, 1983. [Online]. Available: https://tools.ietf.org/rfc/rfc854.txt.

[20] T. Ylonen and C. Lonvick, "The Secure Shell (SSH) Transport Layer Protocol," IETF, 2006. [Online]. Available: https://tools.ietf.org/html/rfc4253.

[21] HL7, "HL7 Version 2 Product Suite," 2019. [Online]. Available: https://www.hl7.org/implement/standards/product_brief.cfm?product_id=185.

[22] NEMA, "DICOM Standard," [Online]. Available: https://www.dicomstandard.org/.

[23] CLSI, "POCT01," [Online]. Available: https://clsi.org/standards/products/point-of-care-testing/documents/poct01/.

[24] CLSI, "LIS02," [Online]. Available: https://clsi.org/standards/products/automation-and-informatics/documents/lis02/.

[25] C. Bodungen, B. Singer, A. Shbeeb, K. Wilhoit and S. Hilt, Hacking Exposed Industrial Control Systems: ICS and SCADA Security Secrets & Solutions, McGraw-Hill, 2016.

[26] P. Ciholas, A. Lennie, P. Sadigova and J. Such, "The Security of Smart Buildings: a Systematic Literature Review," 2019. [Online]. Available: https://arxiv.org/abs/1901.05837.

[27] M. Gabriel, A. Noblin, A. Rutherford, A. Walden and K. Cortelyou-Ward, "Data Breach Locations, Types, and Associated Characteristics Among US Hospitals," The American Journal of Managed Care, vol. 24, no. 2, 2018.

[28] E. Snell, "Hospital Data Breaches Most Common, Affect the Most Patients," Health IT Security, 2018. [Online]. Available: https://healthitsecurity.com/news/hospital-data-breaches-most-common-affect-the-most-patients.

[29] J. Kwon and M. Johnson, "THE MARKET EFFECT OF HEALTHCARE SECURITY: DO PATIENTS CARE ABOUT DATA BREACHES?," in Workshop on the Economics of Information Security (WEIS), 2015.

[30] Symantec, "New Orangeworm attack group targets the healthcare sector in the U.S., Europe, and Asia," 2018. [Online]. Available: https://symantec-enterprise-blogs.security.com/blogs/threat-intelligence/orangeworm-targets-healthcare-us-europe-asia.

[31] FireEye, "DoubleDragon - APT41, a dual espionage and cyber crime operation," [Online]. Available: https://content.fireeye.com/apt-41/rpt-apt41.

[32] S. Mansfield-Devine, "Ransomware: taking businesses hostage," Network Security, 2016.

[33] D. Gayle, A. Topping, I. Sample, S. Marsh and V. Dodd, "NHS seeks to recover from global cyber-attack as security concerns resurface," The Guardian, 2017. [Online]. Available: https://www.theguardian.com/society/2017/may/12/hospitals-across-england-hit-by-large-scale-cyber-attack?CMP=share_btn_tw.

[34] T. Brewster, Forbes, 2017. [Online]. Available: https://www.forbes.com/sites/thomasbrewster/2017/05/17/wannacry-ransomware-hit-real-medical-devices/#25d50bd425cf.

[35] M. Field, "WannaCry cyber attack cost the NHS £92m as 19,000 appointments cancelled," The Telegraph, 2018. [Online]. Available: https://www.telegraph.co.uk/technology/2018/10/11/wannacry-cyber-attack-cost-nhs-92m-19000-appointments-cancelled/.

[36] B. Krebs, "Europe's Largest Private Hospital Operator Fresenius Hit by Ransomware," 2020. [Online]. Available: https://krebsonsecurity.com/2020/05/europes-largest-private-hospital-operator-fresenius-hit-by-ransomware/.

[37] M. Ahmed and A. Ullah, "False data injection attacks in healthcare," in Australasian Conference on Data Mining, 2018.

[38] D. Regalado, "Inside the Alaris Infusion Pump, not too much medicine, plz," DEF CON 25 IoT Village, 2017. [Online]. Available: https://youtu.be/w4sChnS4DrI.

[39] B. Rios, "Infusion Pump Teardown," S4x16, 2016. [Online]. Available: https://youtu.be/pq9sCaoBVOw.

[40] Y. Xu, D. Tran, Y. Tian and H. Alemzadeh, "Analysis of Cyber-Security Vulnerabilities of Interconnected Medical Devices," in IEEE/ACM International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE), 2019.

[41] A. Duggal, "Understanding HL7 2.X Standards, Pen Testing, and Defending HL7 2.X Messages," Black Hat US, 2016. [Online]. Available: https://youtu.be/MR7cH44fjrc.

[42] D. Haselhorst, "HL7 Data Interfaces in Medical Environments: Understanding the Fundamental Flaw in Healthcare," SANS, 2017. [Online]. Available: https://www.sans.org/reading-room/whitepapers/vpns/paper/38005.

[43] D. Haselhorst, "HL7 Data Interfaces in Medical Environments: Attacking and Defending the Achille's Heel of Healthcare," SANS, 2017. [Online]. Available: https://www.sans.org/reading-room/whitepapers/vpns/paper/38010.

[44] M. Bland, C. Dameff and J. Tully, "Pestilential Protocol: How Unsecure HL-7 Messages Threaten Patient Lives," Black Hat US, 2018. [Online]. Available: https://i.blackhat.com/us-18/Thu-August-9/us-18-Dameff-Pestilential-Protocol-How-Unsecure-HL7-Messages_Threaten-Patient-Lives.pdf.

[45] Y. Mirsky, T. Mahler, I. Shelef and Y. Elovici, "CT-GAN: Malicious Tampering of 3D Medical Imagery using Deep Learning," in USENIX Security, 2019.

[46] Cylera, "HIPAA-Protected Malware? Exploiting DICOM Flaw to Embed Malware in CT/MRI Imagery," 2019. [Online]. Available: https://labs.cylera.com/2019/04/16/pe-dicom-medical-malware/.

[47] D. McKee, "80 to 0 in Under 5 Seconds: Falsifying a Medical Patient's Vitals," McAfee, 2018. [Online]. Available: https://securingtomorrow.mcafee.com/other-blogs/mcafee-labs/80-to-0-inunder-5-seconds-falsifying-a-medical-patients-vitals/.

[48] Siemens, DCA Vantage™ Analyzer Host Computer Communications Link, 2011.

[49] Philips, "Data Export Interface Programming Guide," [Online]. Available: http://incenter.medical.philips.com/doclib/enc/fetch/applibid1.DAD/2000/4504/577242/577243/577247/582636/582882/X2%2c_MP%2c_MX_%26_FM_Series_Rel._L.0_Data_Export_Interface_Program._Guide_4535_645_88011_(ENG).pdf%3fnodeid%3d11407611%26vernum%3d-2.

[50] Stilgherrian, "Medical device cybersecurity will be rubbish for 20 more years," ZDNet, 2019. [Online]. Available: https://www.zdnet.com/article/medical-device-cybersecurity-will-be-rubbish-for-20-more-years/.

[51] Forreseter Research, "Strategic Plan: The Zero Trust Security Playbook," 2019. [Online]. Available: https://reprints.forrester.com/#/assets/2/1757/RES137210/reports.

[52] Center for Internet Security, "The 20 CIS Controls & Resources," [Online]. Available: https://www.cisecurity.org/controls/cis-controls-list/.

[53] B. Osborn, J. McWilliams, B. Beyer and M. Saltonstall, "BeyondCorp: Design to Deployment at Google," 2016. [Online].

[54] VMWare, "Introducing VMware Zero Trust Model," 2019. [Online]. Available: https://techzone.vmware.com/blog/introducing-vmware-zero-trust-model.

[55] Microsoft, "Zero Trust Maturity Model," [Online]. Available: https://download.microsoft.com/download/f/9/2/f92129bc-0d6e-4b8e-a47b-288432bae68e/Zero_Trust_Vision_Paper_Final%2010.28.pdf.

ABOUT FORESCOUT

Forescout is the leader in Enterprise of Things security, offering a holistic platform that continuously identifies, segments and enforces compliance of every connected thing across any heterogeneous network. The Forescout platform is the most widely deployed, scalable, enterprise-class solution for agentless device visibility and control. It deploys quickly on your existing infrastructure–without requiring agents, upgrades or 802.1X authentication. Fortune 1000 companies and government organizations trust Forescout to reduce the risk of business disruption from security incidents or breaches, ensure and demonstrate security compliance and increase security operations productivity.

Don't just see it. Secure it.™

Contact us today to actively defend your Enterprise of Things.

FORESCOUT RESEARCH LABS

forescout.com/industries/healthcare

salesdev@forescout.com

toll free 1-866-377-8771

FORESCOUT
Active Defense for the Enterprise of Things™

Forescout Technologies, Inc.
190 W Tasman Dr.
San Jose, CA 95134 USA

Toll-Free (US) 1-866-377-8771
Tel (Intl) +1-408-213-3191
Support +1-708-237-6591

Learn more at Forescout.com