A CYBER RISK SCORING SYSTEM FOR MEDICAL DEVICES

THESIS

Ian W. Stine, Capt, USAF

AFIT-ENG-MS-17-M-072

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

AFIT-ENG-MS-17-M-072

A CYBER RISK SCORING SYSTEM FOR MEDICAL DEVICES

THESIS

Presented to the Faculty

Department of Electrical and Computer Engineering

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

in Partial Fulfillment of the Requirements for the

Degree of Master of Science in Cyber Operations

Ian W. Stine, B.S.C.E.

Capt, USAF

March 2017

AFIT-ENG-MS-17-M-072

A CYBER RISK SCORING SYSTEM FOR MEDICAL DEVICES

THESIS

Ian W. Stine, B.S.C.E.
Capt, USAF

Committee Membership:

LTC Mason J. Rice, PhD
Chair

Lt Col John M. Pecarina, PhD
Member

Stephen J. Dunlap
Member

AFIT-ENG-MS-17-M-072

Abstract

The increased connectivity of medical devices has expedited patient treatment and provides lifesaving capabilities, but a lack of emphasis on device security has led to cybersecurity breaches for many healthcare organizations. Most medical professionals do not have a background in information technology or cybersecurity, yet they are responsible for assessing which treatment provides the best balance of risk and probability for success. This paper presents a two-part risk assessment framework that uses a doctor's worst case assessment of a device's potential to impact a patient and a security questionnaire based on the STRIDE model to generate a risk score on a scale from 0 to 10. Four test cases based on relevant medical devices are used to demonstrate the practical application of the framework.

Table of Contents

Page

List of Figures

List of Tables

Table Page

A CYBER RISK SCORING SYSTEM FOR MEDICAL DEVICES

I. Introduction

1.1 Background

In 1961, Dr. Lawrence L. Weed began developing the first electronic medical record system called the Problem-Oriented Medical Information System (Promise) [14]. His design was intended to reduce paperwork and simplify patient record keeping. Around this same time, Physicians were also developing the first implantable pacemaker [1]. These two electronic systems would eventually be incorporated into the Internet of Things (IoT). Along with record keeping and pacemakers, all manner of devices would eventually be brought online (e.g., thermometers, infusion pumps and electrocardiograms).

In 2009, the US Congress published the Health Information Technology for Economic and Clinical Health (HITECH) Act. This act provided incentives for using (and penalties for not using) Electronic Medical Records (EMR), Electronic Health Records (EHR) or Electronic Patient Records (EPR) [33] and took the first step in a lengthy endeavour to address the need to protect Protected Health Information (PHI) [22]. Over the next four years healthcare organizations began to increase their acceptance of network enabled medical devices and the healthcare sector saw an increasing number of network breaches.

According to Identity Theft Resource Center's Data Breach Reports, from 2009 to 2012 the healthcare industry moved from the fourth most breached industry sector to the second most, trailing only the business sector [22]. This drastic increase in security

breaches occurred immediately after the publication of the HITECH Act. Fortunately, this dramatic increase also raised awareness of the growing problems associated with networked medical devices. The Department of Homeland Security (DHS) issued the National Infrastructure Protection Plan which assigned the responsibility for ensuring the security of public health and health care to the Department of Health and Human Services (DHHS) [16]. Emphasizing the need to address the growing security issues, Presidential Policy Directive 21 (PPD21) was issued in 2013. In the directive, roles and responsibilities, strategic imperatives and implementation guidance were outlined and handed off to sector-specific agencies and once again the DHHS was assigned the responsibility of securing healthcare and public health [19].

Over the past few years, several medical devices have made headlines as independent security researchers found vulnerabilities which allow unauthorized access to the devices or cause the devices to function in a manner outside the manufacturer's original intent. In June 2015, Computerworld [21] published an article which highlighted security firm TrapX's report of hackers using medical equipment as pivot points to gain access to healthcare networks. The article specifically mentions cases where X-ray machines, picture archive and communications systems (PACS) and blood gas analyzers were used for this purpose [21]. Currently, there are no published reports of medical devices being exploited via cyberattacks that resulted in patient harm or death, but that does not mean it isn't possible. Several researchers have focused on cyber-physical effects and found that devices such as infusion pumps, pacemakers and insulin pumps could be manipulated to harm patients [10, 11, 32]. With evidence highlighting the still growing cybersecurity issues industry and governmental agencies recognize the need to improve security in medical networks and devices.

While government and industry work to improve device security, healthcare organizations are left to manage the modernization of their facilities without a complete

understanding of the risks. Cybersecurity risk management needs to be incorporated into the process to ensure patient safety.

1.2 Research Questions

The scoring system presented in this paper is designed to improve the cyber risk assessment process within the healthcare industry for devices that leverage network connections. This scoring system should aid organizations in identifying devices with the potential to harm a patient or greatly affect patient care. To achieve this lofty goal, three key objectives are addressed: (i) ease of use; (ii) low cost; and (iii) easily understood results.

i. Ensure ease of use.

This scoring system is intended to assist medical and management personnel with understanding cyber risk in medical devices. Many of these people may not have formal cybersecurity training. Therefore, the scoring system must be easily understood by someone without a strong cybersecurity background.

ii. Provide a low cost solution.

Many healthcare organizations are already dealing with limited resources so it's important to make sure that use of the scoring system isn't cost prohibitive. This can be obtained by minimizing: (i) equipment costs; (ii) training; and (iii) personnel requirements.

iii. Produce easily understood results.

One hurdle to wide spread adoption is whether or not the system produces usable results. Any additional work that must be done to process and interpret the re-

sults may reduce the odds that it will be incorporated into existing risk management practices. Using a widely adopted industry standard scale with predefined categories will reduce the likelihood that the risk scoring system will produce counter-intuitive results.

1.3 Research Contributions

The primary goal of this research is to design a framework for ranking the risk posed by the cyber-physical interactions of medical devices. Furthermore, this research intends to create a scoring system which can be used as part of a larger cybersecurity risk management framework. The framework and scoring system both support the Department of DHS's efforts to strengthen critical infrastructure.

1.4 Limitations

The test scenarios used during the evaluation process have several assumptions and limitations. Each scenario is based on a medical device which has been produced and marketed to healthcare organizations. The scenarios are framed by environmental and impact assumptions to define each devices' potential impact. Some of the devices have already been examined by cybersecurity experts and their findings have been publicized. In most cases specific device flaws, techniques used and device security features were not published. The security features for each device were determined based on reported vulnerabilities and available device marketing documentation.

1.5 Thesis Layout

Chapter II discusses current healthcare industry regulations, cybersecurity improvement efforts in healthcare and relevant works used to develop the scoring system. Chapter III details the design of the risk assessment framework and the scoring

system. Chapter IV presents the results and findings of the scenarios. Chapter V analyzes the results from the scenarios, presents future work and summarizes the thesis topic.

II. Related Research

2.1 Regulation of Medical Device Cybersecurity

47 years after the first implantable pacemaker, the Joint Commission on Accreditation of Healthcare Organizations released a Sentinel Event Alert regarding the safe implementation of health information and converging technologies. In the alert, the commission highlighted many factors which contributed to harmful computer technology-related errors. Factors included an over-reliance on vendor advice (without the oversight of an objective third party), not carefully considering the impact technology can have on care processes and the failure to quickly fix technology when it becomes counterproductive [13].

The production and marketing of medical devices is highly regulated. For a device to legally reach market it must conform to specific requirements outlined in the Code of Federal Regulations (CFR) and pass through an approval process conducted by the FDA.

Code of Federal Regulations.

The CFR is the codification of the general and permanent rules published in the Federal Register by the departments and agencies of the Federal Government [31]. Title 21 of the CFR contains the parts which are regulated by DHHS and the FDA. Part 800 - Part 898 specifically address the manufacturing and security of medical devices. Within the CFR, emphasis is placed on ensuring patient and operator safety. As part of that effort, one of the key regulatory requirements is that a device be classified as either Class 1, Class 2 or Class 3 based on its safety and clinical effectiveness. All of the regulations use broad terms to discuss the documents and content which the manufacturer must present for a medical device to gain approval. Interpretation and

enforcement of these requirements is left up to the agency responsible for enforcing the CFR.

Food and Drug Administration Guidance.

The FDA is a subordinate organization to DHHS. With regards to medical devices, the FDA openly claims responsibility for: (i) ensuring the safety, efficacy and security of medical devices; and (ii) advancing the public health by helping speed innovations that make medical products more effective, safer and more affordable [30]. The connection to the regulatory process made the FDA uniquely positioned to influence and regulate medical device cybersecurity. To enhance industry understanding of their interpretation of Title 21 of the CFR, the FDA periodically releases guidance which represents their current thinking on a specific topic [29].

In January 2005, the FDA released the Guidance for Industry - Cybersecurity for Networked Medical Devices Containing Off-the-Shelf (OTS) Software [26]. In this document, the FDA identifies the manufacturer as the responsible party for continued safe and effective performance of the medical device, including the performance of OTS software that is part of the device. This document also addresses reporting requirements for cybersecurity patches by stating that in most cases manufacturers would not need to report a cybersecurity patch so long as they have evaluated the change and recorded the correction in the product's records. However, if the software patch impacts the safety or effectiveness of the medical device, manufacturers should report the correction to the FDA [26].

While the OTS guidance specifically states that it is intended for devices containing OTS software, these concepts are incorporated on a larger scale into the two most recent guidance documents: (i) Content of Premarket Submissions for Management of Medical Devices (issued October 2, 2014); and (ii) the draft Postmarket

Management of Medical Devices (issued January 22, 2016). The premarket guidance provides recommendations and information which manufacturers should include in FDA medical device premarket submissions for effective cybersecurity management and is intended to reduce the risk to patients by decreasing the likelihood that device functionality is intentionally or unintentionally compromised by inadequate cybersecurity [24]. The FDA guidance ties suggestions for validation and risk analysis back to 21 CFR 820.30(g) which is the requirement for manufacturers to conduct design validation. This linkage allows the FDA to reject submissions that have not complied with requirements. There are five elements that the FDA recommends for cybersecurity and vulnerability management approaches [24]: (i) identification of assets, threats and vulnerabilities; (ii) assessment of the impact of threats and vulnerabilities on device functionality and end users/patients; (iii) assessment of the likelihood of a threat and of a vulnerability being exploited; (iv) determination of risk levels and suitable mitigation strategies; and (v) assessment of residual risk and risk acceptance criteria. Each of these elements expects manufacturers to self-identify the vulnerabilities, associated risks for their device and provide actions taken to secure the identified issues.

Where the premarket guidance focused on requirements to receive device approval, the postmarket guidance clarifies life cycle management recommendations and emphasizes that manufacturers should monitor, identify and address cybersecurity vulnerabilities and exploits [28]. The guidance also clarifies their expectations with regards to the reporting requirements. For the majority of cases, actions taken by manufacturers to address cybersecurity vulnerabilities and exploits are considered routine updates or patches, which do not require advance notification or reporting under 21 CFR part 806. For a small subset of cybersecurity vulnerabilities that may compromise the essential clinical performance of a device (and present a reasonable probability

of serious adverse health consequences or death), the FDA requires medical device manufacturers to notify the Agency [28].

The postmarket guidance also references Executive Order (EO) 13691 - Promoting Private Sector Cybersecurity Information Sharing. EO 13691 encourages the development of Information Sharing Analysis Organizations (ISAOs) to serve as focal points for information sharing and collaboration within the private sector and between the private sector and government. In addition to encouraging information sharing, ISAOs collect data and participate in collaborative establishment of standards and best practices relating to cybersecurity [23]. EO 13691 outlines measures that should be taken to ensure collected information is treated as Protected Critical Infrastructure Information [28]. The post market guidance continues by outlining key definitions and elements which should be incorporated as part of any postmarket risk management strategy. Most of these elements are consistent with those identified in the NIST Framework for Improving Critical Infrastructure Cybersecurity which can be found in Table 1 on page 14.

All of the FDA's guidance documents present well-supported recommendations for improving device cybersecurity. While each document is representative of the FDA's views, there is a statement included with each guidance document that may undermine many of their suggestions: "FDA's guidance documents, including this draft guidance, do not establish legally enforceable responsibilities. Instead, guidances describe the Agency's current thinking on a topic and should be viewed only as recommendations, unless specific regulatory or statutory requirements are cited. The use of the word should in Agency guidance means that something is suggested or recommended, but not required" [28]. The flexibility of the FDA's guidance allows manufacturers to make their own decisions about cybersecurity. However, the lack of concrete legal requirements makes enforcing cybersecurity requirements difficult.

2.2 Current State of Cybersecurity in Healthcare

From a cybersecurity perspective, the primary organizations reporting on safety and security for medical devices are the Industrial Controls Systems Cyber Emergency Response Team (ICS-CERT) and the FDA. While ICS-CERT has published numerous alerts and advisories warning about vulnerable devices, the FDA has rarely officially recalled a device due to cyber vulnerabilities [27]. The FDA did however, release an unprecedented alert for the Hospira Symbique Infusion Pumps which were identified as having a vulnerability which could potentially allow an attacker to control the dosage of the medication to the patient. This indicates the FDA intends to treat cyber vulnerabilities in a manner similar to other device flaws [9]. Identification of high risk devices could enable the FDA to outline increased testing requirements for these devices and provide a means to categorize devices based on their potential cybersecurity risk.

Risk Framework Implementation Trends.

The healthcare industry appears to be adapting as cybersecurity awareness increases. In March 2016, Dimensional Research released the results of a survey which focused on the adoption of security frameworks [5]. In the survey, 338 participants were asked a wide range of questions to understand which security frameworks were adopted, motivations for adoption and how fully they were adopted. In these results, adoption of at least one security framework was the norm, however, healthcare showed the lowest adoption percentage with only 61% stating that they adopted a risk management framework. This falls short of even the education industry which reported 77% adoption and even lower than the 83-88% reported by the other four industries (i.e., banking, IT, government and manufacturing). The survey also reported that only 3% of the respondents used a security framework other than the top four: (i)

Payment Card Industry Data Security Council Standard (PCI); (ii) NIST framework; (iii) CIS Critical Security Controls; or (iv) ISO/IEC 27001/27002. While the NIST framework was reported as being the lowest of the top four, the report indicated that by the end of 2016 its adoption rate should grow from 29% to 43%. This indicates that even though the NIST framework is relatively new, many consumers are accepting the NIST framework. Of the participants who reported adopting the NIST framework, more than half indicated that a significant investment is needed to fully conform with the five functions [5]. From the implementation levels of other industries it is clear that healthcare is behind and is in need of widespread implementation.

Mayo Clinic.

In 2013, the Mayo Clinic invited computer security experts to test the security of approximately 40 medical devices [20]. Every device tested revealed substantial flaws. One of the hackers noted that in comparison to other sensitive industries, hospitals appeared to be roughly a decade behind the standard security curve. As a result of these tests the Mayo Clinic began updating their IT security practices. Some of the changes included adding language to documents which specifically addresses cybersecurity in medical devices and adding to their in-house security team. One of the documents updated was their information technology medical equipment proposal questionnaire, which asks various questions about the product's capabilities, life cycle and account management and whether the manufacturer will allow independent testing [3, 15]. Security experts applauded the Mayo Clinic's efforts but recognize that few hospitals have the resources (or influence) to implement these security measures [20].

Improving the Current State.

Most of the recent improvements focus on addressing the lack of baked-in security in medical devices and improving the communication within the industry. Both of these efforts will take time to fully mature and see measurable results. In the meantime, using and improving risk management practices within organizations may help protect patients and healthcare organizations from vulnerable devices. As discussed earlier, one possible reason robust risk management practices aren't being adopted is the investment requirements. The Mayo Clinic had the resources to establish their own internal IT security group in an attempt to identify flawed devices and conduct risk management processes. For organizations that lack the resources necessary to implement a security team, a low cost, easy to use system for identifying and ranking higher risk devices would allow incorporation of new or improved risk management practices into the decision making process.

2.3 Foundational Components

With the added emphasis being placed on securing healthcare networks, it is inevitable that industry partners and regulatory agencies will begin to insight change. These changes are intended to promote better awareness of cybersecurity and improve the security of these devices and networks. The risk scoring system presented in this paper is intended to be used as part of a risk management framework (e.g., NIST Security Framework) and was designed using current industry concepts presented by the Escal Institute of Advanced Technologies (SANS Institute) and the Forum of Incident Response and Security Teams (FIRST).

NIST Security Framework.

The NIST framework was developed in response to EO 13636 which called for a voluntary risk-based cybersecurity framework to help organizations manage cybersecurity risks. The framework is designed to be used by multiple organizations responsible for securing critical infrastructure and is currently being endorsed by the FDA [24, 28]. To account for variations in these systems, the framework contains modular components which can be implemented based on the desired risk management process results [17].

The NIST framework establishes a hierarchical structure by using five core functions to organize basic cybersecurity activities. The core functions and their descriptions are in Table 1. Each function is then divided into categories which are based on cybersecurity outcomes tied to programmatic needs and particular activities. Subcategories within each category represent specific outcomes of technical or management activities. While the subcategories are not exhaustive, they provide a set of results that help support the outcomes in each category. The framework ties informative references to each subcategory in an attempt to provide framework users with standards, guidelines and practices common among critical infrastructure sectors [17]. The flexible nature of the NIST framework makes it a great choice for healthcare organizations, but full implementation could require more resources than what are available. The risk scoring system presented in this paper is designed to help alleviate some of the resource requirements by simplifying the process of calculating and ranking cybersecurity risk for medical devices.

FIRST Common Vulnerability Scoring System.

The Common Vulnerability Scoring System (CVSS), created by the Forum of Incident Response and Security Teams (FIRST), is designed to be a robust scoring

Table 1. NIST framework core functions.

Core Function	Description
Identify:	Develop the organizational understanding to manage cybersecurity risk to systems, assets, data and capabilities.
Protect:	Develop and implement the appropriate safeguards to ensure delivery of critical infrastructure services.
Detect:	Develop and implement the appropriate activities to identify the occurrence of a cybersecurity event.
Respond:	Develop and implement the appropriate action regarding a detected cybersecurity event.
Recover:	Develop and implement the appropriate activities to maintain plans for resilience and to restore any capabilities that were impaired due to a cybersecurity event.

system for IT vulnerabilities [6]. The scoring system was developed using collaborative input and is free to use. Several entities including the National Vulnerability Database (NVD) and the Common Vulnerabilities and Exposure (CVE) use the scoring system to enhance their databases [18].

CVSS version 3 generates three metric groups: (i) Base; (ii) Temporal; and (iii) Environmental. The Base group represents intrinsic qualities about the vulnerability, the Temporal group reflects the characteristics of a vulnerability that change over time and the Environmental group represents the characteristics of a vulnerability that are unique to a user's environment. The Base metrics produce a score ranging from 0.0 to 10.0, which can then be modified by scoring the Temporal and Environmental metrics [6]. The CVSS calculator created by FIRST presents the user with the various metrics divided into the three metric groups with a series of defined values for each metric. Once a value is selected for each metric within a specific group, the data is combined into a vector string and used to generate the score(s). The Base score must be generated to receive a Temporal and Environmental score [8]. Each of these scores combines more and more information to provide a more accurate score for the

vulnerability based on a given situation. These scores can then be used as part of a larger risk management process. Overall, the design of the calculator allows for a simplistic user interface requiring minimal knowledge of the scoring system.

SANS Attacker Objectives.

The Escal Institute of Advanced Technologies (more commonly referred to as the SANS Institute) is a cooperative research and education organization which specializes in information security and cybersecurity. Their InfoSec Reading Room published a paper by Michael Assante and Robert Lee [2] discussing the industrial control systems cyber kill chain. In that paper they define nine attacker objectives based on the types of effects. The categories are represented by three types of effects (i.e., loss, denial and manipulation) and these three effects could be seen in up to four different areas of the system: (i) The view; (ii) The controls; (iii) Safety; and (iv) Sensors. The categories and areas are presented in Table 2. A modified version of these categories is used by the risk scoring system presented in this paper to describe the potential affects that could be seen in a medical device.

Table 2. SANS attacker objectives.

Action	System Component
Loss of:	View Control
Denial of:	View Control Safety
Manipulation of:	View Control Sensors Safety

III. Cyber Risk Scoring System

3.1 Risk Scoring System

With the various accounts of vulnerable devices reaching market, swift action is needed to help identify which devices pose a greater risk to patient safety. This framework is designed to aid in the identification of devices which have the potential to harm a patient or greatly affect a physician's ability to provide an accurate diagnosis. To achieve this goal there are three key objectives that need to be addressed: (i) ease of use; (ii) low cost; and (iii) easily understood results.

Ease of Use.

The acquisition process varies greatly within the healthcare industry. In many cases, purchases are made through Group Purchasing Organizations or by groups within a single organization. In either case, the people deciding which devices to purchase lack the necessary skills to identify and rank cybersecurity risks. To maximize the framework's usability the inputs need to be easy to understand. To accomplish this, all inputs are presented in a questionnaire format (i.e., yes/no or column of values) similar to the method used by the CVSS.

Low Cost.

The cost of using this framework is relatively low when compared to creating an in-house cybersecurity team. All that is required is someone with medical knowledge who can accurately answer the outcome assessment questions and someone who can input the answers to the security questions presented as part of the device assessment. Understandably, healthcare organizations may not know the answers to the device assessment portion but the manufacturer should be able to provide those answers via

a questionnaire similar to that used by the Mayo Clinic in their risk management process [15].

Understandable Results.

Arguably the most important part of any scoring system is how easily the user can interpret the results. To make this process as easy as possible, the framework was built so that the results align with the existing vulnerability severity ratings outlined by the NVD. These ratings are also being used to describe the values output by CVSS. As mentioned before, CVSS is used to score existing vulnerabilities on a scale of zero to ten (ten being the worst). The NVD rating system further describes these values by creating three distinct categories: (i) Low: 0.0-3.9; (ii) Medium: 4.0-6.9; and (iii) High: 7.0-10.0 [18].

3.2 Framework

This scoring system is built upon a two part framework: (i) a worst-case assessment of a device's potential outcome if it were to be compromised; and (ii) an assessment of the device's security features. These two components provide the inputs for the scoring system. Future work is need to incorporate environmental security factors into the scoring system.

Outcome Assessment.

The FDA's postmarket guidance recognizes the difficult nature of estimating the probability of a cybersecurity exploit occurring. In the absence of sufficient probability data, the FDA suggests using a reasonable worst-case estimate or setting the default value of the probability to one [28]. When the probability of an event occurring is set to one, this represents the inevitability that the event will happen.

Following this suggestion, the outcome assessment of this framework is built on the assumption that the device will be compromised and that the worst-case outcome should be used as the basis for the scoring system. Since this framework is assessing the outcome of events that haven't yet occurred, there must be description of what events are possible.

The attacker objectives created by the SANS Institute form a solid basis for describing potential events and can be found in Table 2. When looking at the subcategories in relation to medical devices there are a lot of similarities. All medical devices have some form of human machine interface (HMI), a mechanism for controlling the device and safety features to ensure user and patient safety. In most situations, external sensors are directly connected to or co-located with the medical device. In order for the sensor to be manipulated, physical modification would be required or the data would need to be modified in transit. Identifying this limitation allows for the removal of the attacker objective which describes actions affecting sensors (i.e., manipulation of sensors) and the manipulation of the data in transit can be accounted for as a manipulation of the view. To simplify the assessment process further, limiting the number of categories with minimal differences was important. Since the end goal is to assess the device's potential to impact a patient or diagnosis, loss or denial of any of the subcategories will often represent a similar end-state. Therefore, the various subcategories were combined and the final list of possible effects is listed in Table 3. These five effects are the situations for which the device's outcome assessment will be made.

The final piece is to define what possible outcomes could occur in each of the five attacker objective categories. The FDA's postmarket guidance references ANSI/AAMI/ISO 14971:2007/(R)2010: Medical Devices–Application of Risk Management to Medical Devices which provides five key categories to which the FDA provided

Table 3. Possible cyber effects.

Possible Effects
Loss or Denial of View
Loss or Denial of Control
Manipulation of View
Manipulation of Control
Denial or Manipulation of Safety

possible descriptions as seen in Table 4 [28].

Table 4. Severity levels.

Common Term	Description
Negligible:	Inconvenience or temporary discomfort.
Minor:	Results in temporary injury or impairment not requiring professional medical intervention.
Serious:	Results in injury or impairment requiring professional medical intervention.
Critical:	Results in permanent impairment or life-threatening injury.
Catastrophic:	Results in patient death.

While the Severity levels describe the potential effects on a patient there is no mention of the impact the device may have on a diagnosis. To do this the potential impacts must be categorized. One simple solution is to look at whether the device can affect a diagnosis and then determine if the device in question is the sole source of the information, if a secondary device is presenting redundant data or if data is corroborated by additional data. By following this logical progression and using the same terminology suggested by the FDA, four potential descriptions are created. The catastrophic outcome is not definable for impacting the diagnosis since it already results in patient death as part of the original definition. The descriptions of the effects on the diagnosis can be found in Table 5.

Now that there are descriptions for the effect on the patient and on the diagno-

Table 5. Device potential diagnosis outcomes.

Common Term	Description
Negligible:	Inconvenience to the staff with no effect on the diagnosis.
Minor:	Low potential for misdiagnosis with additional redundant data sources available.
Serious:	Potential for misdiagnosis with additional collaborative data sources available.
Critical:	Results in misdiagnosis with no additional independent data sources.
Catastrophic:	N/A.

sis the descriptions can be merged into a single set. This new set will be used for categorizing the device's worst-case potential outcome for each of the attacker objectives. Table 6 contains the final descriptions for each of the common terms used in the device assessment. Note that catastrophic refers to direct patient death and does not contain a description of the impact on the diagnosis. This was done intentionally since direct loss of life presents a higher risk to patient than indirect effects which may be prevented through other environmental variables. Table 7 is an example of the the assessment decision matrix. The first five rows contain the attacker objectives and will receive a single mark in the column representing the worst potential outcome for each event. The last row is used to display the summation of the totals in each column. This information will later be used by the scoring system to calculate an appropriate risk score.

Device Assessment.

The device assessment is the second half of the risk scoring framework and is designed to assess whether the manufacturer addressed basic cybersecurity concepts when producing the device. Microsoft developed the STRIDE model for classifying threats. Each threat is associated with a specific IT property which helps combat the

Table 6. Device potential outcomes.

Common Term	Description
Negligible:	Inconvenience or temporary discomfort or inconvenience to the staff with no effect on the patient.
Minor:	Results in temporary injury or impairment not requiring professional medical intervention or low potential for misdiagnosis with additional redundant data sources available.
Serious:	Results in injury or impairment requiring professional medical intervention or potential for misdiagnosis with additional collaborative data sources available.
Critical:	Results in permanent impairment or life-threatening injury or results in misdiagnosis with no additional independent data sources.
Catastrophic:	Results in patient death.

Table 7. Device assessment table.

Attacker Objective	Negligible	Minor	Serious	Critical	Catastrophic
Loss/Denial of View					
Loss/Denial of Control					
Manipulation of View					
Manipulation of Control					
Denial or Manipulation of Safety					
Column Totals					

threat [4]. Table 8 shows these relationships.

Microsoft also provides standard mitigation techniques which can be used to support each of these areas [4]. Using their suggestions as a foundation, a set of questions were created which target each of the six properties associated with the STRIDE

Table 8. STRIDE model.

Threat	Property
Spoofing	Authentication
Tampering	Integrity
Repudiation	Non-repudiation
Information Disclosure	Confidentiality
Denial of Service	Availability
Elevation of Privilege	Authorization

model. The questions and associated properties can be seen in Table 9.

Table 9. STRIDE properties with associated device security questions.

Property	Security Questions
Authentication:	Does the system use multi-factor authentication? Does the system enforce secure credential creation, usage and maintenance principles?
Integrity:	Can the system detect and prevent manipulated parameters? Does the system protect against tampering and reverse engineering? Were secure software design principles followed during development (to include third party software)?
Non-repudiation:	Does the system verify and log all user actions with attribution?
Confidentiality:	Does the system follow industry standard encryption practices to secure connections?
Availability:	Was the system built and tested for high availability (e.g., fuzz testing and load testing)?
Authorization:	Does the system allow for management of all users and privileges?

Scoring System.

The FDA is currently recommending CVSS to categorize vulnerabilities in medical devices. CVSS uses a 0–10 scale whose scores are broken into the severity levels outlined by NVD [28]. Creating a scoring system that ties into this existing scale

ensures that users of both systems will be able to recognize the risk relationship associated with devices. In addition to using a 0–10 scale, CVSS version 3.0 also generates three separate scores. Each score takes into account additional factors to produce a progressively more accurate and valuable score [7]. Following this strategy, the outcome assessment is used to generate a basic score which simply describes the devices potential to impact a patient's health or the doctor's ability to provide an accurate diagnosis.

Since the NVD severity levels have three key areas, the first step was to define what meaning, if any, should be tied to these boundaries. Since minimizing patient death is the top concern, devices with possible outcomes containing values scored as critical or catastrophic should all end up in the highest of the three ranges (7–10). It was also determined that a 2:1 ratio would be reasonable for each of the levels, meaning two critical ratings should represent a similar value to one catastrophic, two serious ratings should be scored similarly to one critical and so on. With the desire to restrict catastrophic outcomes to the upper score range and the relationship of values established, it seemed reasonable to build the scoring system around three key ideas: (i) All devices with at least one catastrophic outcome should have a base score greater than or equal to 7; (ii) A ratio of 2:1 should be maintained between each category (e.g., a device with two critical values should be roughly equivalent to a device with one catastrophic value); and (iii) No device can score higher than a 10.0.

If the values of each category were set as constants, the range of values would be difficult to fit on a 0–10 scale without modification (e.g., if the catastrophic outcome was set to 7, to ensure the minimum threshold was met, a device with two or more catastrophic values would already exceed the scales maximum value of 10). To solve this problem and fulfil the three desired parameters, a scaling value system based on the number of critical and catastrophic outcomes was developed. The scoring

algorithm looks at the inputs and determines their individual values based on the number of values contained in the critical or catastrophic categories. Table 10 contains the individual values of each category based on the total number of potential outcomes rated as critical or catastrophic. Figure 1 shows the ranges of potential values based on the number of critical or catastrophic outcomes. Table 11 identifies upper and lower limits within each range depending on the presence (or absence) of catastrophic outcomes. The base score is a summation of the determined worst-case value in each of the categories and can be seen in Equation 1.

$$BaseScore = \Sigma[AttackerObjectiveScore] \tag{1}$$

Table 10. Attacker objective score based on outcomes.

Number of Critical/Cat	Negligible	Minor	Serious	Critical	Catastrophic
0	0	.3	1	N/A	N/A
1	0	.25	.4	3.5	6
2	0	.2	.3	3	3.5
3	0	.1	.2	2.1	2.8
4	0	.05	.1	1.8	2.3
5	N/A	N/A	N/A	1.5	2

Now that the base score is established from the outcome assessment, this score needs to be augmented by the results of the device assessment to achieve the device's risk score. As mentioned earlier, each question is associated with a specific property for which it is designed to assess. These properties have the potential to impact multiple attacker objectives that are used for the outcome assessment but might have no impact on other objectives. The association matrix in Table 12 displays

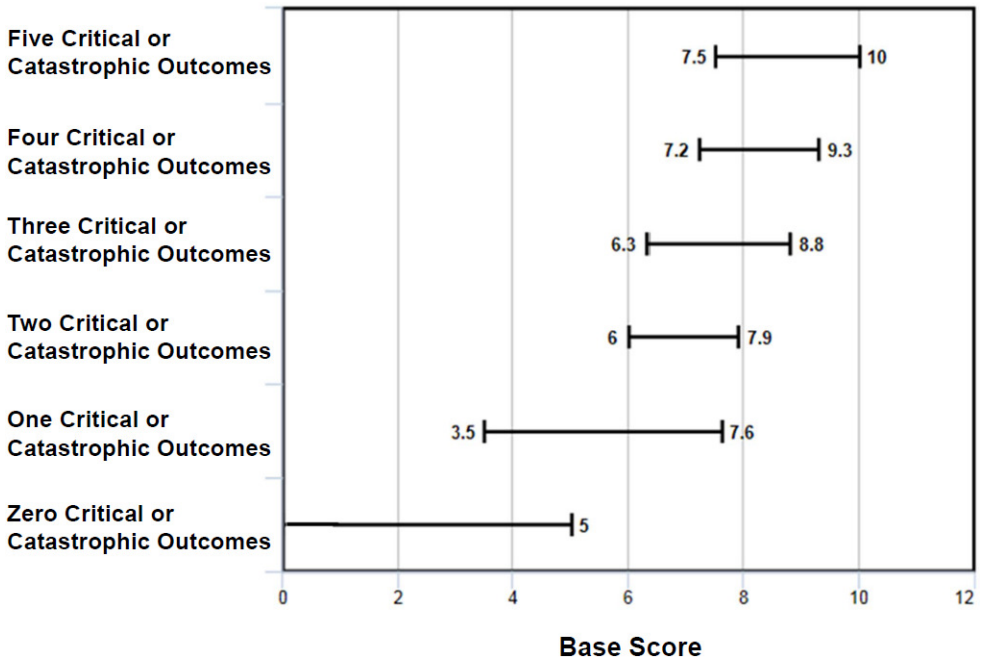

Figure 1. Ranges of potential values based on the number of critical and catastrophic outcomes.

Table 11. Key potential values based on the number of critical and catastrophic outcomes.

Number of Critical/Cat	Lowest Value without a Cat	Lowest Value with a Cat	Highest Value without a Cat	Highest Possible Value
0	0	N/A	5	5
1	3.5	6	5.1	7.6
2	6	6.5	6.9	7.9
3	6.3	7	6.7	8.8
4	7.2	7.7	7.3	9.3
5	7.5	8	7.5	10

how the various security principles (i.e., authentication, integrity, non-repudiation, confidentiality, availability and authorization) are tied to the attacker objectives. Each attacker objective that is marked as being impacted by a security principle can receive adjustment when the associated security question is answered. The association matrix will add some weight to the categories, but overall their individual values

will be represented consistently (i.e., implementing security measures that address all questions within a specific property will affect the score equally given that their impact within the association matrix is identical).

Table 12. Property association matrix.

Attacker Objective	Authen.	Integrity	Non Rep.	Conf.	Avail.	Authorization
Loss/Denial of View	Yes	Yes	No	No	Yes	No
Loss/Denial of Control	Yes	Yes	No	No	Yes	No
Manipulation of View	Yes	Yes	Yes	Yes	No	Yes
Manipulation of Control	Yes	Yes	Yes	Yes	No	Yes
Denial or Manipulation of Safety	Yes	Yes	Yes	Yes	No	Yes

Some properties have multiple questions associated with them. In these cases, each question is weighted equally and the proportional value will impact the score (e.g., a device meeting 2 out of 3 of the security questions in integrity will receive 2/3 credit as opposed to a full 1 credit for meeting all 3). The final algorithms used to determine the device's risk score are Equation 2 and Equation 3. In the Equation 2, the OldAttackerObjectiveScore is the value associated with each individual attacker objective which were used in the summation to calculate the base score. The assessment credit is divided by 20 to provide reasonable impact scaling. Note that CVSS also uses constant values in their equations [7].

$$NewAttackerObjectiveScore = OldAttackerObjectiveScore \\ * (1 - \Sigma[DeviceAssessmentCredit]/20) \tag{2}$$

$$DeviceRiskScore = \Sigma[NewAttackerObjectiveScores] \tag{3}$$

IV. Test Scenarios

The following three scenarios involve three real devices that have documented vulnerabilities. In each of the three cases the specifics of any exploits were kept confidential, but the potential to cause harm to patients were made public. Since the CVSS scoring system presents all scores rounded to the tenths position, all calculations presented in the scenarios are rounded in a similar manner.

4.1 Digital Thermometer

This scenario involves a device that is used to determine a patient's temperature external to the body. The device displays the reading on an LCD and transmits the data via Bluetooth connection to a paired device. Since this device does not have a direct method of impacting the patient through normal operations, the potential impact on the patient should fall into the negligible categories and in the worst case, a physician might need to confirm inaccurate readings with a second device or may simply touch the patient's forehead to validate readings. These effects would clearly dictate the negligible category be used for all five of the attacker objectives. This would result in an outcome assessment score of 0 for the device. Since the first stage of the assessment scores a 0, the second stage would also result in a 0 regardless of which security measures were taken. These scores represent an innocuous device that has no potential to directly harm a patient or produce a misdiagnosis.

4.2 Medication Delivery Device

This scenario involves a device that controls medication being delivered to a patient. The device has networking capabilities and at least one remote connection. Since this is an assessment of the worst case potential outcome, this scenario assumes

that the medication being delivered is either lifesaving or has the potential to be lethal with improper dosage. It also assumes that the delivery schedule must be maintained or a critical level event could occur. The outcome assessment input can be viewed in Table 13. In this case there are 5 critical or catastrophic values. Using the information provided by the outcome assessment the values for each attacker objective can be obtained from Table 10. 13 also provides the individual attacker objective values in the summation of the individual columns. Inserting the appropriate values into Equation 1 produces an outcome assessment score of 9 which represents one of the highest possible scores and means that this device may need additional testing to ensure security measures have been properly implemented.

$$\begin{aligned} BaseScore &= \Sigma[AttackerObjectiveScore] \\ &= 1.5 + 1.5 + 2 + 2 + 2 = 9 \end{aligned}$$

Table 13. Medication delivery device outcome assessment.

Attacker Objective	**Negligible**	**Minor**	**Serious**	**Critical**	**Catastrophic**
Loss or Denial of View				X	
Loss or Denial of Control				X	
Manipulation of View					X
Manipulation of Control					X
Denial or Manipulation of Safety					X
Column Totals				1.5+1.5=3	2+2+2=6

The second half of the assessment requires the security features be analyszed using the questionnaire. The contents of Table 14 were used for the device assessment.

Using the responses in the device assessment the property association matrix in Table 15 can be filled in with the assessed values. These values are the individual device assessment credits which are needed in Equation 2. Table 16 shows the original and final component scores as well as the calculations. The resulting risk assessment score of 7.6 was calculated for the medication delivery device. The minimum score achievable based on the outcome assessment is 7.1, meaning the device covered several security properties but didn't cover them all.

Table 14. Medication delivery device security questionnaire.

Property	Security Questions	Yes/No
Authentication:	Does the system use multifactor authentication?	No
	Does the system enforce secure credential creation, usage and maintenance principals?	No
Integrity:	Can the system detect and prevent manipulated parameters?	No
	Does the system protect against tampering and reverse engineering?	Yes
	Were secure software design principles followed during development of all software (to include thrid party software)?	Yes
Non Repudiation:	Does the system verify and log all user actions with attribution?	Yes
Confidentiality:	Does the system follow industry standard encryption practices to secure all connections?	Yes
Availability:	Was the system built and tested for high availability (e.g., fuzz testing and load testing)?	Yes
Authorization:	Does the system allow for management of all users and privileges?	Yes

The Hospira infusion pumps were found to contain flaws which could allow an unauthorized user to control the device and change the dosage the pump delivers [25]. This test scenario is similar to what might have been uncovered in the risk assessment process of the Hospira pumps. The scores calculated in each step of the

Table 15. Medication delivery device value adjustment matrix.

Attacker Objective	Authen.	Integrity	Non Rep.	Conf.	Avail.	Authorization
Loss/Denial of View	0	2/3	N/A	N/A	1	N/A
Loss/Denial of Control	0	2/3	N/A	N/A	1	N/A
Manipulation of View	0	2/3	1	1	N/A	1
Manipulation of Control	0	2/3	1	1	N/A	1
Denial or Manipulation of Safety	0	2/3	1	1	N/A	1

Table 16. Medication delivery device attacker objective risk score components.

Attacker Objective	Old Component Score	Equation	New Component Score
Loss/Denial of View	1.5	1.5 * (1 - (0+2/3+0+0+1+0)/20	1.4
Loss/Denial of Control	1.5	1.5 * (1 - (0+2/3+0+0+1+0)/20	1.4
Manipulation of View	2	2 * (1 - (0+2/3+1+1+0+1)/20	1.6
Manipulation of Control	2	2 * (1 - (0+2/3+1+1+0+1)/20	1.6
Denial or Manipulation of Safety	2	2 * (1 - (0+2/3+1+1+0+1)/20	1.6

process had the potential to raise awareness of the device's potential to harm patients and could have been used to prompt further testing to ensure the device was properly secured.

4.3 Implantable Device

This scenario involves a device that is surgically implanted into a patient. Once implanted, the device is controlled via a wireless connection and is used to regulate the patient's heart. If the device becomes unresponsive to controls, a critical level event could occur (e.g., the patient would require surgery to replace the device). If the controlling application fails to display device information a serious level event could occur (e.g., misdiagnosis of current health condition due to the device being the sole source data). Manipulation of the device outside the set parameters, or failure of the device's safety controls could result in a catastrophic event (e.g., patient dies due to device hyperactivity). The outcome assessment for this scenario can be seen in Table 17. Since there are 4 critical or catastrophic values, Table 10 can be used to identify the individual values. Based on these values, a potential impact score of 8.8 was calculated using Equation 1. This score falls into the high risk category and represents a device that may need additional validation of its security measures.

$$\begin{aligned} BaseScore &= \Sigma[AttackerObjectiveScore] \\ &= .1 + 1.8 + 2.3 + 2.3 + 2.3 = 8.8 \end{aligned}$$

The second half of the assessment requires the security features be analyzed using the questionnaire. The contents of Table 18 were used for the device assessment and the property association matrix in Table 19 was filled in with the resulting values. Table 20 shows the original and final component values and the use of Equation 2. The summation of these values generated a risk score of 8.8. This score is representative of a manufacturer who didn't take security into account when developing their product and is a clear warning to potential users that the device is likely exploitable.

The results of this scenario are based around the findings of Daniel Halperin et al.

Table 17. Implantable device outcome assessment.

Attacker Objective	Negligible	Minor	Serious	Critical	Catastrophic
Loss or Denial of View			X		
Loss or Denial of Control				X	
Manipulation of View					X
Manipulation of Control					X
Denial or Manipulation of Safety					X
Column Totals			.1	1.8	2.3+2.3+2.3=6.9

[10] and Barnaby Jack [12]. Their research into implantable cardioverter-defibrillators (ICDs) revealed that the devices could be reprogrammed to shock the patient. These devices were equipped with RF transmitters that contained vulnerabilities that were intended to allow doctors to interact with the ICD without having to operate on the patient [10, 12].

4.4 Electrocardiogram

This scenario uses an electrocardiogram (ECG) which connects via a Bluetooth connection to a laptop (or smart device) which is running a patient monitor application. While the device does monitor a vital process in the human body, it's intended for use in situations where patient health is not in immediate danger (e.g., diagnostic testing). This limiting factor dictates that the biggest impact the device could have on the patient is through manipulation of the physician's diagnosis by presenting false readings. In situations where this is the case, there will often be other devices with

Table 18. Implantable device security questionnaire.

Property	Security Questions	Yes/No
Authentication:	Does the system use multifactor authentication?	No
	Does the system enforce secure credential creation, usage and maintenance principals?	No
Integrity:	Can the system detect and prevent manipulated parameters?	No
	Does the system protect against tampering and reverse engineering?	No
	Were secure software design principles followed during development of all software (to include third party software)?	No
Non Repudiation:	Does the system verify and log all user actions with attribution?	No
Confidentiality:	Does the system follow industry standard encryption practices to secure all connections?	No
Availability:	Was the system built and tested for high availability (e.g., fuzz testing and load testing)?	No
Authorization:	Does the system allow for management of all users and privileges?	No

Table 19. Implantable device value adjustment matrix.

Attacker Objective	Authen.	Integrity	Non Rep.	Conf.	Avail.	Authorization
Loss/Denial of View	0	0	N/A	N/A	0	N/A
Loss/Denial of Control	0	0	N/A	N/A	0	N/A
Manipulation of View	0	0	0	0	N/A	0
Manipulation of Control	0	0	0	0	N/A	0
Denial or Manipulation of Safety	0	0	0	0	N/A	0

Table 20. Implantable device attacker objective risk score components.

Attacker Objective	Old Component Score	Equation	New Component Score
Loss/Denial of View	.1	.1 * (1 - (0+0+0+0+0+0)/20	.1
Loss/Denial of Control	1.8	1.8 * (1 - (0+0+0+0+0+0)/20	1.8
Manipulation of View	2.3	2.3 * (1 - (0+0+0+0+0+0)/20	2.3
Manipulation of Control	2.3	2.3 * (1 - (0+0+0+0+0+0)/20	2.3
Denial or Manipulation of Safety	2.3	2.3 * (1 - (0+0+0+0+0+0)/20	2.3

the potential to contradict the data presented by the device. This may limit the worst case potential for any manipulation of the view or controls to the serious category. If alternate methods where available to verify the readings, the worst case potential would be lowered to the minor category. Since the device does not have the capacity to physically impact the patient, manipulation or loss of safety has a negligible outcome. The loss or denial of the view or controls would also be negligible given the original assumption that the device is only being used to take readings during non life-threatening situations. Given these assessments of the worst case potential outcomes, Table 10 can be used to identify the associated values for the case where there are no critical or catastrophic ratings. The resulting assessment and values can be seen in Table 21. Using these values with Equation 1 results in a base score of 2. This represents a device which possesses little or no capability to directly harm a patient and with careful use would rarely affect a diagnosis. However, if this device were to be used outside of the defined parameters (e.g., as a monitor during surgery or in an emergency situation) the worst case outcomes would be drastically different.

$$BaseScore = \Sigma[AttackerObjectiveScore]$$

$$= 0 + 0 + 1 + 1 + 0 = 2$$

Table 21. Electrocardiogram device outcome assessment for intended use.

Attacker Objective	**Negligible**	**Minor**	**Serious**	**Critical**	**Catastrophic**
Loss or Denial of View	X				
Loss or Denial of Control	X				
Manipulation of View			X		
Manipulation of Control			X		
Denial or Manipulation of Safety	X				
Column Totals	0+0+0=0		1+1=2		

The device security questionnaire for the ECG scenario can be seen in Table 22. Next, the responses from the questionnaire are entered into the value adjustment matrix in Table 23 and the device adjustment values are calculated. These values are then entered into Equation 2 as seen in Table 24. Summation of these values results in a risk score of 2.0 which represents a lack of security being implemented in the device. (Note that without rounding the individual attacker objective scores, the risk score would be 1.9 instead of 2.0). Since this device was intended for non-emergency situations the lack of security doesn't present as big a threat to patient health. If this device were to be used in an environment where it was the sole source of lifesaving information to medical personnel any loss or manipulation of the device's view or

controls could lead to catastrophic outcomes. Table 25 contains the assessment if the device were used in this type of situation (e.g., an ambulance). Application of these new values to Equation 1 results in a new base score of 9.2. This score is drastically different from the earlier assessment due to the patient's condition and lack of redundancy. Since the device security assessment remains the same, the same adjustment matrix can be used and the calculations can be seen in Table 26. The lack of security results in a risk score of 9.2. (Again, note that without rounding the individual attacker objective scores, the risk score would be 9.0 instead of 9.2). This is an extremely high risk score and indicates that alternatives should be considered.

As highlighted by this example, a device can receive drastically different scores depending on the operational environment in which it will be used.

Table 22. Electrocardiogram device security questionnaire.

Property	Security Questions	Yes/No
Authentication:	Does the system use multifactor authentication?	No
	Does the system enforce secure credential creation, usage and maintenance principals?	No
Integrity:	Can the system detect and prevent manipulated parameters?	No
	Does the system protect against tampering and reverse engineering?	Yes
	Were secure software design principles followed during development of all software (to include third party software)?	No
Non Repudiation:	Does the system verify and log all user actions with attribution?	No
Confidentiality:	Does the system follow industry standard encryption practices to secure all connections?	No
Availability:	Was the system built and tested for high availability (e.g., fuzz testing and load testing)?	No
Authorization:	Does the system allow for management of all users and privileges?	No

$$BaseScore = \Sigma[AttackerObjectiveScore]$$
$$= 2.3 + 2.3 + 2.3 + 2.3 + 0 = 9.2$$

Table 23. Electrocardiogram device value adjustment matrix.

Attacker Objective	Authen.	Integrity	Non Rep.	Conf.	Avail.	Authorization
Loss/Denial of View	0	1/3	N/A	N/A	0	N/A
Loss/Denial of Control	0	1/3	N/A	N/A	0	N/A
Manipulation of View	0	1/3	0	0	N/A	0
Manipulation of Control	0	1/3	0	0	N/A	0
Denial or Manipulation of Safety	0	1/3	0	0	N/A	0

As with the other scenarios, the ECG scenario was based off a real device, the Nasiff CardioCard and was investigated as part of this research effort. During the testing process it was discovered that the device contained a default static pairing key that could not be managed. This could allow a Bluetooth enabled device within range to execute a denial of service on any active device. With the static nature of the device and the lack of security, it is possible that the ECG could be vulnerable to spoofing or a man-in-the-middle attack which could result in the manipulation of vital readings.

Table 24. Electrocardiogram attacker objective risk score components for intended use.

Attacker Objective	Old Component Score	Equation	New Component Score
Loss/Denial of View	0	0 * (1 - (0+1/3+0+0+0+0)/20	0
Loss/Denial of Control	0	0 * (1 - (0+1/3+0+0+0+0)/20	0
Manipulation of View	1	1 * (1 - (0+1/3+0+0+0+0)/20	1
Manipulation of Control	1	1 * (1 - (0+1/3+0+0+0+0)/20	1
Denial or Manipulation of Safety	0	0 * (1 - (0+1/3+0+0+0+0)/20	0

Table 25. Electrocardiogram device outcome assessment for unintended use.

Attacker Objective	Negligible	Minor	Serious	Critical	Catastrophic
Loss or Denial of View					X
Loss or Denial of Control					X
Manipulation of View					X
Manipulation of Control					X
Denial or Manipulation of Safety	X				
Column Totals	0				2.3+2.3+2.3+2.3=9.2

Table 26. Electrocardiogram device attacker objective risk score components for unintended use.

Attacker Objective	Old Component Score	Equation	New Component Score
Loss/Denial of View	2.3	2.3 * (1 - (0+1/3+0+0+0+0)/20	2.3
Loss/Denial of Control	2.3	2.3 * (1 - (0+1/3+0+0+0+0)/20	2.3
Manipulation of View	2.3	2.3 * (1 - (0+1/3+0+0+0+0)/20	2.3
Manipulation of Control	2.3	2.3 * (1 - (0+1/3+0+0+0+0)/20	2.3
Denial or Manipulation of Safety	0	0 * (1 - (0+1/3+0+0+0+0)/20	0

V. Recommendations and Conclusions

5.1 Analysis of Results and Future Work

The four scenarios presented in this paper represent devices on the market which have identified vulnerabilities. Applying the risk scoring framework to each of the devices resulted in scores representative of their potential to harm a patient. The scoring system identifies devices with potential outcomes in the catastrophic categories, except in cases where all of the following criteria are met: (i) there are two or fewer attacker objectives rated as critical or catastrophic; (ii) only a single attacker objective receives a catastrophic rating; and (iii) at least one of the remaining objectives is deemed to be negligible. When these specific criteria are met, the resulting score is still greater than or equal to 6 which is still relatively high on the overall scale. Attempts were made to raise these lower values into the desired range but the artificial inflation yielded scores which were not representative of their relative severity (i.e., devices with a single catastrophic outcome could score higher than a device with two). To keep the relative score balance it was deemed acceptable to have outliers containing a single catastrophic outcome which fall below the critical category score threshold of 7.

The scoring system presented in this paper provides a solution for one of the key categories in risk management process (i.e., risk assessment). The simple interface increases the system's ease of use and limits the training needed to use it. There is little or no cost associated with implementing the scoring system which makes it accessible to even the most resource restricted organizations. Moreover, it uses a scale which is already widely used and provides a method of pro-actively identifying medical devices which represent greater cybersecurity risk.

5.2 Future Work

Feedback and Scoring Data.

Presenting this risk scoring system to physicians could generate invaluable feedback and additional scoring data. Feedback regarding the system's usefulness in active risk management processes could potentially highlight areas for improvement. Additional feedback regarding ease of use, processes replaced by system and suggestions regarding ways to improve the system would also be beneficial.

Expansion of Risk Factors.

It would also be beneficial if the risk scoring system could account for environmental factors (e.g., network security features). Similar scoring systems have followed a similar advancement process where over time additional factors were added to the algorithms to account for the complex nature of a networked environment. Expanding the assessment to account for other devices on the network (e.g., vulnerable devices or high risk devices on the network) and standard network security practices (e.g., network segregation and isolation) would greatly increase the systems ability to assess potential risk for specific network environments.

Acceptance and Publication.

For this risk scoring system to be widely adopted it needs support from larger organizations. NIST and the FDA both supported CVSS and helped increase its popularity. Getting these two organizations to promote the system by including it in their recommendations and guidance documents would increase the system's chances of becoming an industry standard.

Additional Applications.

While the initial intent of this research was to improve the cyber risk management process of healthcare organizations, it is possible that there may be other applications for this type of information. Since the FDA is required to rank devices based on the device's capabilities, it might be useful if the base score of the device impacted the classification and auditory processes. This would help raise FDA awareness for devices that pose high risk to patient health if subjected to cyber attacks. Raising awareness of potential issues before they reach market would help better protect healcare organizations and the patients they are treating.

Another possible application would be in the cyber insurance assessment process. Companies offering cyber insurance could provide incentives to organizations who are actively managing their cyber risk or to those organizations who are allowing patients to have input into the type and amount of cyber risk they are being exposed to. Being able to present applicable data about the types of cyber risks your organization is taking could provide insight into the management processes being used by the healthcare organizations and allow for better management of the financial investment required to maintain adequate cyber insurance.

5.3 Conclusions

Over the last few years, a collective effort has been made to identify ways to increase cybersecurity within the healthcare industry. The results of the collaboration have produced risk management frameworks and processes designed to protect the industry and patients. The risk management practices are extensive and the perceived costs are high. With limited resources many organizations are slow to adopt complex standardized cybersecurity risk management processes.

The proposed risk scoring system, which uses a worst-case assessment of the po-

tential outcomes and assesses compliance with basic cybersecurity. It uses a simple interface (similar to CVSS) which is easy to use, has a low operational cost due to minimal training or personnel requirements and produces recognizable results by utilizing a common scoring scale and rating system to provide consistent scoring for medical devices based on their potential to impact patient health. The scenarios represent a sampling of modern medical devices with known vulnerabilities and illustrate the risk scoring framework's ability to provide useful metrics for risk management programs.

Appendix A. Reverse Engineering Nasiff CardioCard

Open Source Information

The first step in Reverse Engineering the Nasiff CardioCard was to search for open source information that might provide insight into the device's functionality. Basic documents regarding functionality, application and features were readily available but no technical documents were found. Focusing the search on FCC documentation revealed that the device was approved for marketing but no review documents were available.

Reverse Engineering the Hardware

Opening up the device revealed that security measures were taken to scrub identifying marks from the various chips.

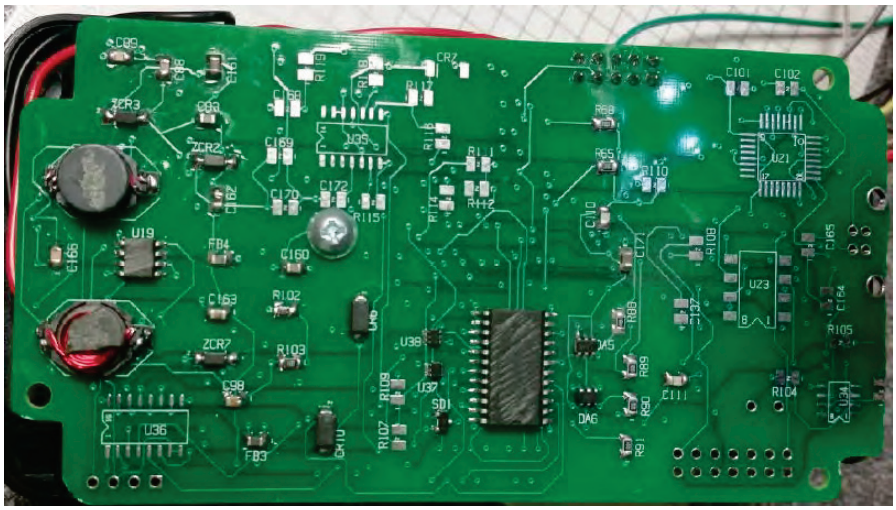

Figure 2. Chip details removed from board 1 side 1.

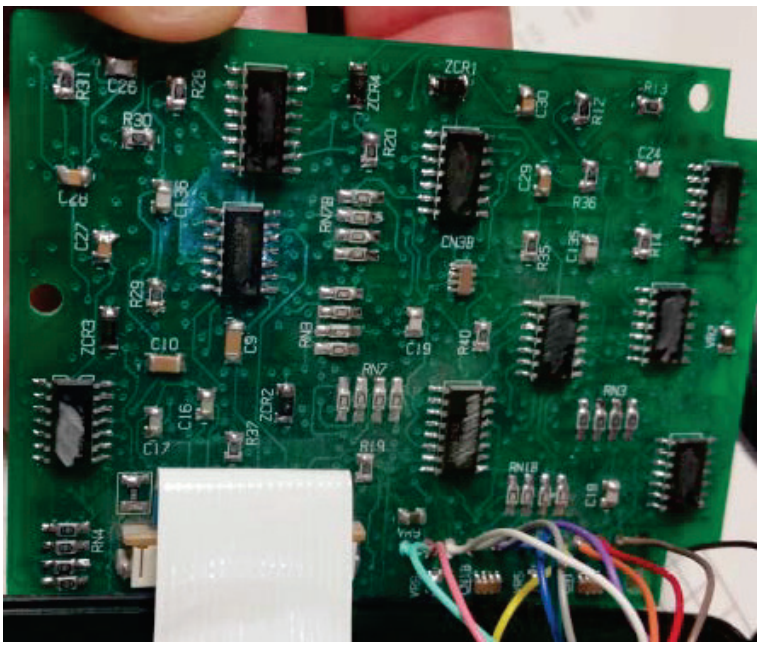

Figure 3. Chip details removed from board 1 side 2.

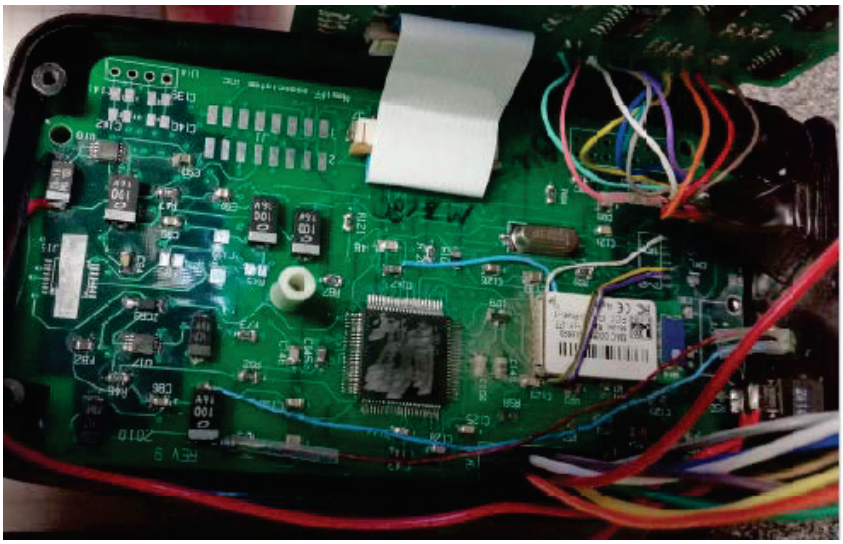

Figure 4. Chip details removed from board 2 side 1.

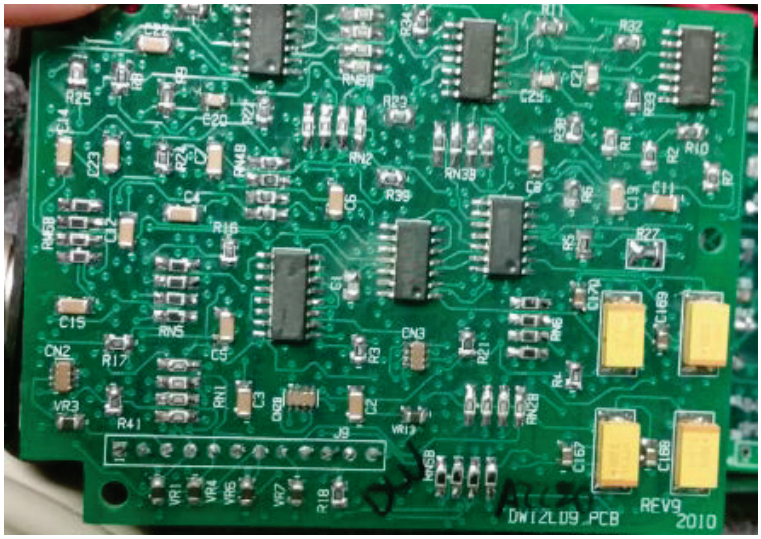

Figure 5. Chip details removed from board 2 side 2.

Using an electronic microscope, several partial numbers were readable.

Figure 6. Chip under USB microscope.

The partial numbers were used to generate possible chip profiles and analysis revealed the presence of the following chips.

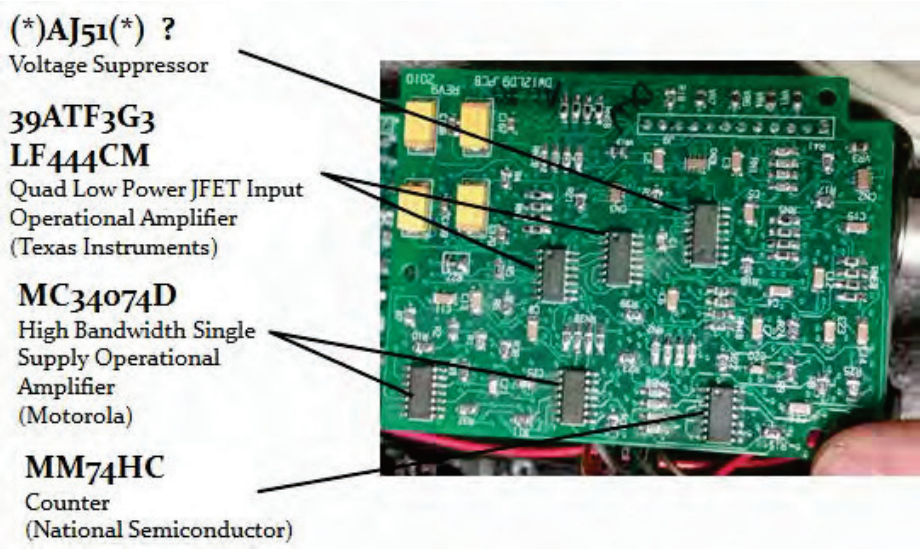

Figure 7. Possible chip profiles for board 1 side 1.

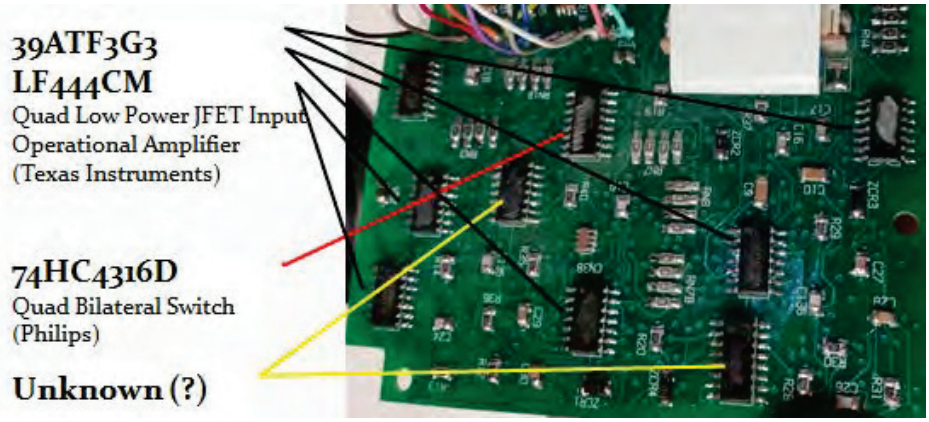

Figure 8. Possible chip profiles for board 1 side 2.

Two of the unidentified chips were located in the same area as amplifiers. It is possible that they are also amplifiers since the system is capable of handling a 12 lead harness and only 10 amplifiers were identifiable.

Inspection of the board also revealed a JTAG port with intact connections located near the microcontroller.

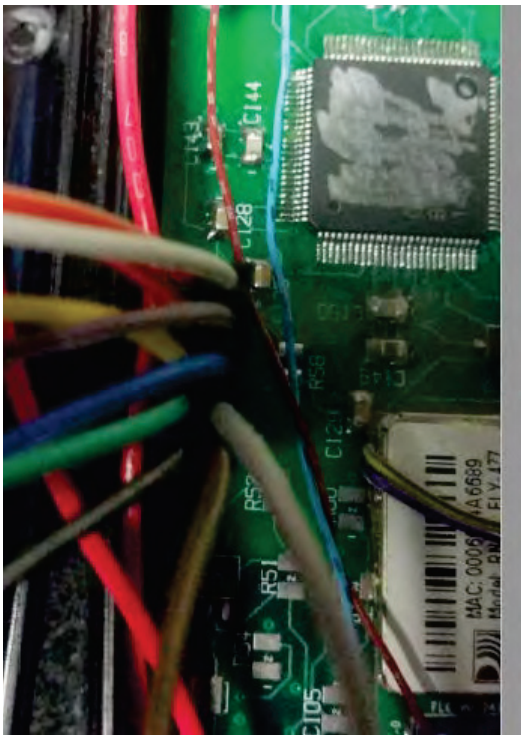

Figure 9. Wired JTAG port.

Using a JTAGULATOR the pinout of the port was identified.

Table 27. Decoded JTAG pinout.

TDO	Yellow Wire
TCK	Green Wire
TDI	Purple Wire
TMS	Blue Wire
Ground	Black Wire/White Wire
Vcc	Red Wire

Once the configuration was determined, the JTAGULATOR was then used to investigate the port and additional information was revealed about the JTAG con-

nection.

Table 28. Device information discovered using JTAGULATOR.

Device ID	E0000803
JEDEC ID	0x401
Part number	0x0000
Version number	0xE

The processor was not clearly identified during inspection. Without knowing the specific instructions needed to recover data from shift registers, additional information was not easily retrieved. If these commands could be determined, it might be possible to retrieve sensitive system data or reprogram the device.

Examining the Communication Link and Protocol

Based on open source documentation the device uses Bluetooth Classic to communicate with the main application. Open source information also revealed that the application needed to access the CardioCard was only available for specific devices (i.e., Microsoft and Apple). The device's power switch is used to activate the device and begin the broadcast sequence announcing its availability to establish pairing. Initial testing revealed that the device could be discovered by an unsupported device using standard operating procedures.

Figure 10. Device located using unsupported Android device.

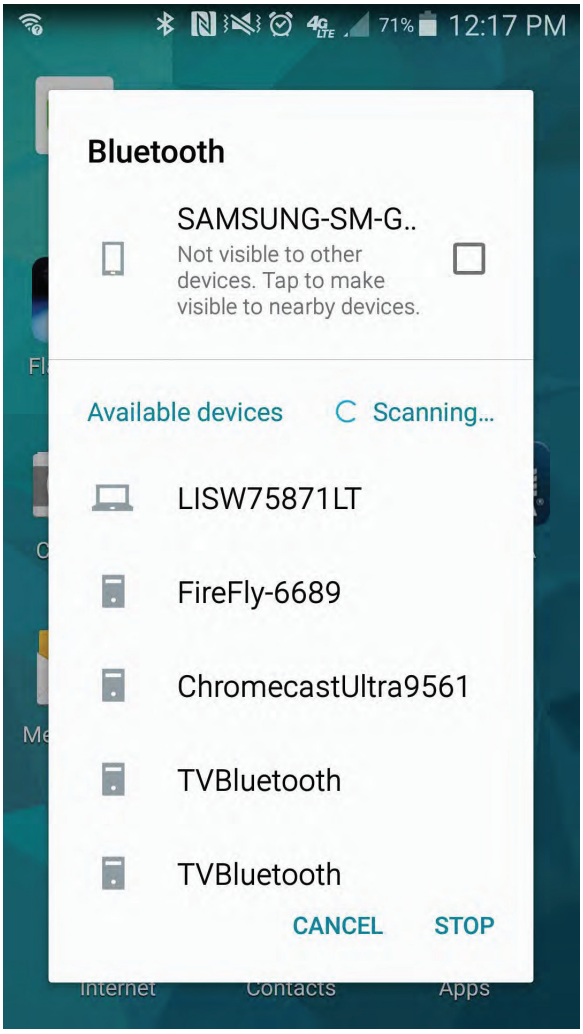

To test these findings an android phone was used to attempt a connection. When connecting with the Android device, a pairing key was requested and two simple pairing keys were suggested by the Android device (0000 and 1234).

Figure 11. Attempting pairing process with unsupported device.

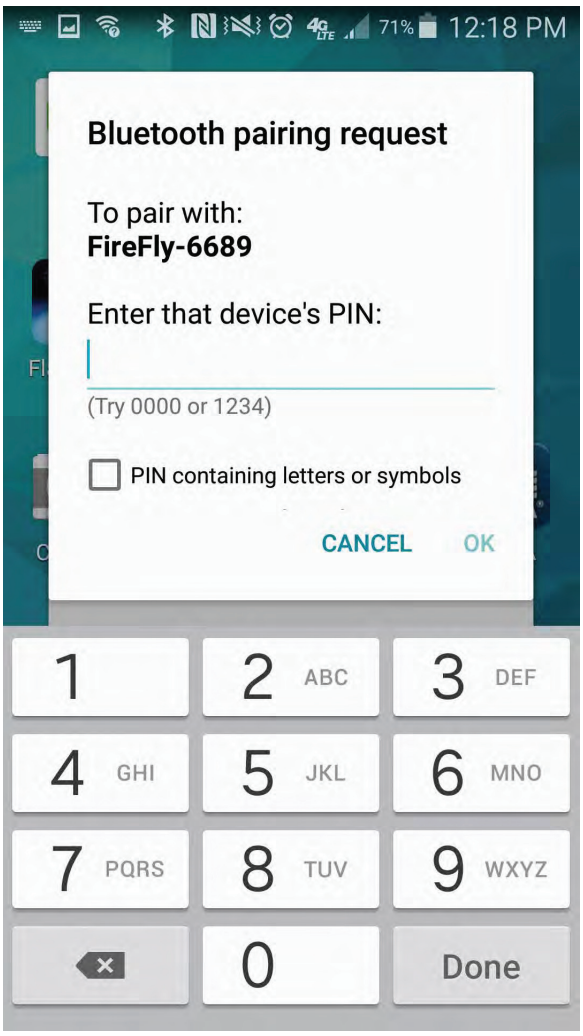

Using the 1234 key, a successful connection was established.

Figure 12. Successful pairing with unsupported device.

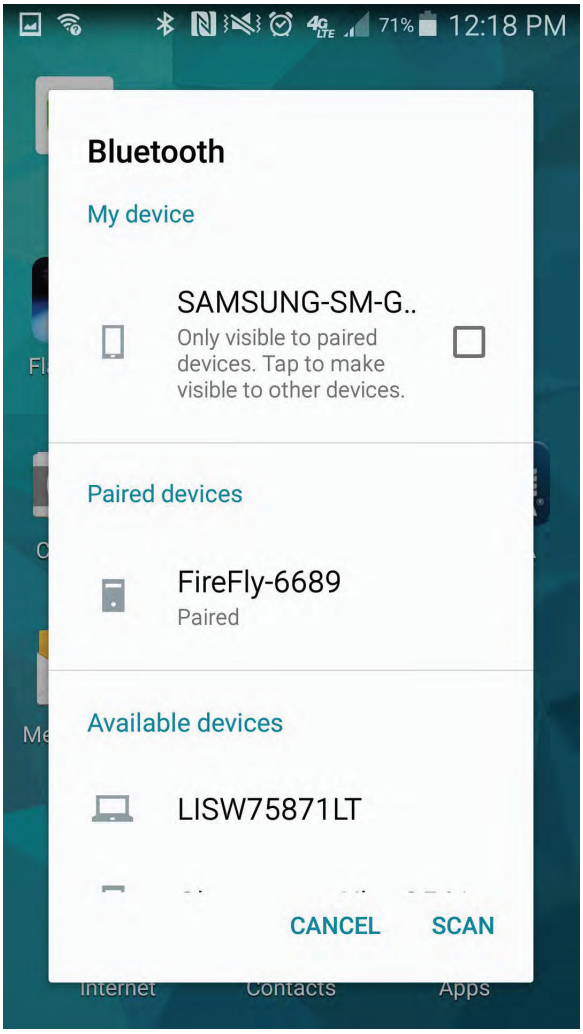

Additional attempts were successful using the same process and key. The repeated success indicates that the CardioCard is using a static pairing key. No additional form of authentication was found during examination of the device.

Inspection of the internal components of the device revealed a Bluetooth module with an FCC identification number.

Figure 13. Identified Bluetooth module with intact FCC ID.

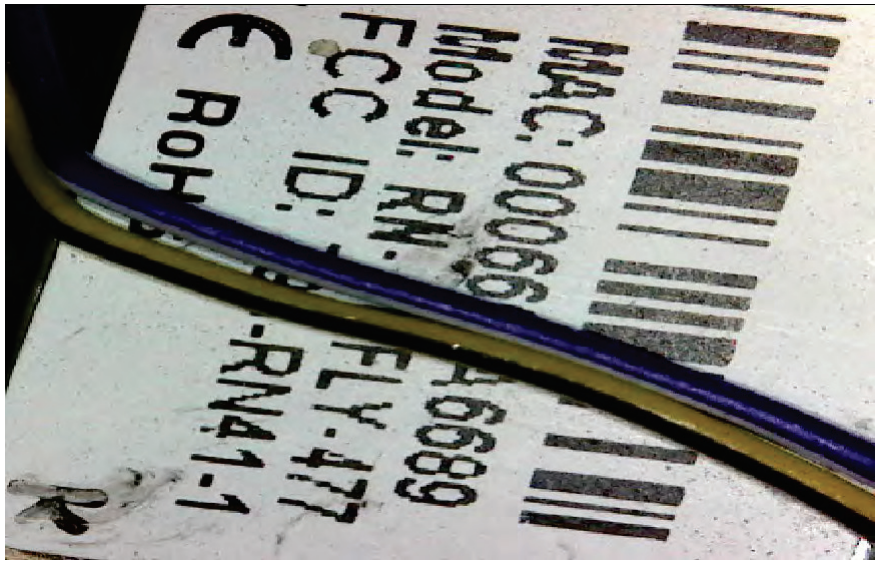

A search of the FCC database produced approval documentation which included design documents and specifications regarding paring capabilities. The documents identified the static pairing key is set by default and configured to be 1234. The module's manufacturer noted that it is possible to change the pairing key if desired. With the static nature of the pairing key and the CardioCard's inability to change the key, anyone with a Bluetooth enabled device could connect with the CardioCard. Since the CardioCard only allows for a single connection to be established at a time, this easily accessed connection has the potential to allow the device to be the target of a Denial of Service Attack.

The next step was to examine the data flowing between the CardioCard and the device hosting the application. An Ubertooth One was used to capture the traffic and packet information was analyzed using the terminal window and Wireshark. The device broadcast its name and was easily located by the Ubertooth One.

Figure 14. Determining UAP/LAP with Ubertooth One and hcitool scan.

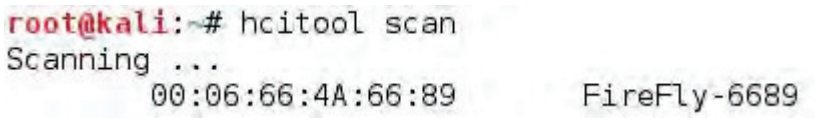

Sniffing traffic allowed the Ubertooth One to identify the upper and lower address parts (UAP and LAP) needed to follow the CardioCard's frequency hopping sequence. Following the CardioCard's sequence allowed the Ubertooth to capture packets being sent by the CardioCard to the application host. By dissecting the traffic, the device hosting the application was identified and then the process was repeated to identify to application host's UAP and LAP. Examination of the traffic revealed a consistent sequence of messages were being generated for command and control.

Summary

One of the basic cybersecurity principles is to ensure passwords are protected and managed appropriately. Based on the FCC documentation, the Bluetooth module being used has a static default pairing key but it can be managed based on implementation. The FDA openly recognizes the security flaws associated with static unmanageable passwords and lists it as one of the programming practices that should be avoided. This design flaw leaves the device wide open for anyone to access and is a gateway to more detrimental exploits. The risk scoring system presented in this paper uses two questions which target static, unmanageable passwords and keys. This information is gathered during the device security assessment and help highlight devices which don't follow basic cybersecurity principles.

Bibliography

1. L. Altman, Arne H. W. Larsson, 86; had first internal pacemaker, *The New York Times*, January 18, 2002.

2. M. Assante and R. Lee, The Industrial Control System Cyber Kill Chain, SANS Institute InfoSec Reading Room, SANS Institute, Bethesda, Maryland (www.sans.org/reading-room/whitepapers/ICS/industrial-control-system-cyber-kill-chain-36297), 2015.

3. Association for the Advancement of Medical Instrumentation News, Mayo Clinic Emphasizes Security with Device Vendors, Arlington, Virginia (www.aami.org/productspublications/articledetail.aspx?ItemNumber=3199), 2016.

4. Department of Electrical Engineering and Computer Science Berkeley, Introduction to Microsoft Software Development Lifecycle (SDL) Threat Modeling, University of California, Berkeley, Berkeley, California (people.eecs.berkeley.edu/~daw/teaching/cs261-f12/hws/Introduction_to_Threat_Modeling.pdf).

5. Dimensional Research, Trends in Security Framework Adoption: A Survey of IT and Security Professionals, San Francisco, California (static.tenable.com/marketing/tenable-csf-report.pdf), 2016.

6. Forum of Incident Response and Security Teams, Common Vulnerability Scoring System, V3 Development Update, Morrisville, North Carolina (www.first.org/cvss), 2015.

7. Forum of Incident Resopnse and Security Teams, Common Vulnerability Scoring System Version 3.0: Specification Document, Morrisville, North Carolina (www.first.org/cvss/specification-document), 2015.

8. Forum of Incident Response and Security Teams, Common Vulnerability Scoring System Version 3.0 Calculator, Morrisville, North Carolina (www.first.org/cvss/calculator/3.0), 2015.

9. P. Fowler, FDA Issues Unprecedented Alert Over Medical Device Cyber Security Risk, Snell & Wilmer, Phoenix, Arizona (www.swlaw.com/blog/product-liability-update/2015/08/17/fda-issues-unprecedented-alert-over-medical-device-cyber-security-risk), 2015.

10. D. Halperin, T. Heydt-Benjamin, B. Ransford, S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno and W. Maisel, Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses, *Proceedings of the IEEE Symposium on Security and Privacy*, pp. 129–142, 2008.

11. Industrial Control Systems Cyber Emergency Response Team, Advisory (ICSA-15-174-01) Hospira Symbiq Infusion System Vulnerability, US Department of Homeland Security, Washington, DC (ics-cert.us-cert.gov/advisories/ICSA-15-174-01), 2015.

12. B. Jack, "Broken Hearts": How Plausible Was the Homeland Pacemaker Hack?, IOActive, Seattle, Washington (blog.ioactive.com/2013/02/broken-hearts-how-plausible-was.html), 2013.

13. Joint Commission on Accreditation of Healthcare Organizations, Safely Implementing Health Information and Converging Technologies, Sentinel Event Alert Issue 42, Oakbrook Terrace, Illonois (www.jointcommission.org/assets/1/18/SEA_42.PDF), 2008.

14. S. Lohr, The 'miracle' of digital health records, 50 years ago, *The New York Times*, 2012.

15. Mayo Clinic, Medical Equipment Procurement Questionnaire, Rochester, Minnesota (www.mayo.edu/pmts/mc7200-mc7299/mc7231.pdf), 2016.

16. J. McNeill and R. Weitz, How to Fix Homeland Security Critical-Infrastructure Protection Plans: A Guide for Congress, The Heritage Foundation, Washington, DC (www.heritage.org/research/reports/2010/04/how-to-fix-homeland-security-critical-infrastructure-protection-plans-a-guide-for-congress), 2010.

17. National Institute of Standards and Technology, Framework for Improving Critical Infrastructure Cybersecurity (Version 1.0), Gaithersburg, Maryland (www.nist.gov/sites/default/files/documents/cyberframework/cybersecurity-framework-021214.pdf), 2014.

18. National Institute of Standards and Technology, NVD Common Vulnerability Scoring System Support v2, Gaithersburg, Maryland (nvd.nist.gov/cvss.cfm), 2014.

19. B. Obama, Presidential Policy Directive – Critical Infrastructure Security and Resilience, PPD-21, The White House, Washington, DC, 2013.

20. M. Reel and J. Robertson, It's way to easy to hack the hospital, *Bloomberg Businessweek* (www.bloomberg.com/features/2015-hospital-hack), 2015.

21. D. Storm, MEDJACK: Hackers hijacking medical devices to create backdoors in hospital networks, *Computerworld* (www.computerworld.com/article/2932371/cybercrime-hacking/medjack-hackers-hijacking-medical-devices-to-create-backdoors-in-hospital-networks.html), 2015.

22. L. Thomson, Health care data breaches and information security, in *Health Care IT: The Essential Lawyers Guide to Health Care Information Technology and the Law*, A. Peabody (Ed.), ABA Book Publishing, Chicago, Illonois, 2013.

23. U.S. Department of Homeland Security, Information Sharing and Analysis Organizations, Washington, DC (www.dhs.gov/isao), 2016.

24. U.S. Food and Drug Administration, Content of Premarket Submissions for Management of Cybersecurity in Medical Devices: Guidance for Industry and Food and Drug Administration Staff, Silver Springs, Maryland (www.fda.gov/ucm/groups/fdagov-public/@fdagov-meddev-gen/documents/document/ucm356190.pdf), 2014.

25. U.S Food and Drug Administration, Cybersecurity Vulnerabilities of Hospira Symbiq Infusion System: FDA Safety Communication, Silver Springs, Maryland (www.fda.gov/MedicalDevices/Safety/AlertsandNotices/ucm456815.htm), 2015.

26. U.S. Food and Drug Administration, Guidance for Industry – Cybersecurity for Networked Medical Devices Containing Off-the-Shelf (OTS) Software, Silver Springs, Maryland (www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm077812.htm), 2016.

27. U.S. Food and Drug Administration, Medical Device Recalls, Silver Springs, Maryland (www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfRES/res.cfm), 2016.

28. U.S Food and Drug Administration, Postmarket Management of Cybersecurity in Medical Devices: Draft Guidance for Industry and Food and Drug Administra-

tion Staff, Silver Springs, Maryland (www.fda.gov/ucm/groups/fdagov-public/@fdagov-meddev-gen/documents/document/ucm482022.pdf), 2016.

29. U.S Food and Drug Administration, Search for FDA Guidance Documents, Silver Springs, Maryland (www.fda.gov/RegulatoryInformation/Guidances/default.htm), 2016.

30. U.S Food and Drug Administration, What We Do, Silver Springs, Maryland (www.fda.gov/AboutFDA/WhatWeDo/default.htm), 2016.

31. U.S. Government, Code of Federal Regulation, Title 21 (Food and Drug), 800 Series, Washingon, DC, 2016.

32. E. Weise, Johnson & Johnson warns of insulin pump hack risk, *USA Today* (www.usatoday.com/story/tech/news/2016/10/04/johnson-johnson-warns-insulin-pump-hack-risk-animas/91542522), October 5, 2016.

33. Wikipedia, Electronic Health Records (en.wikipedia.org/wiki/Electronic_health_record), 2016.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704–0188

1. REPORT DATE *(DD–MM–YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From — To)*
23–03–2017	Master's Thesis	June 2015 — Mar 2017

4. TITLE AND SUBTITLE

A Cyber Risk Scoring System for Medical Devices

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Stine, Ian, W, Capt, USAF

5d. PROJECT NUMBER

17ENG310

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/EN)
2950 Hobson Way
WPAFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT-ENG-MS-17-M-072

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Department of Homeland Security ICS-CERT
POC: Neil Hersheld, DHS ICS-CERT Technical Lead
ATTN: NPPD/CSC/NCSD/US-CERT Mailstop: 0635
245 Murray Lane, SW, Bldg 410, Washington, DC 20528
Email: ics-cert@dhs.gov phone: 1-877-776-7585

10. SPONSOR/MONITOR'S ACRONYM(S)

DHS ICS-CERT

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A:
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The increased connectivity of medical devices has expedited patient treatment and provides lifesaving capabilities, but a lack of emphasis on device security has led to cybersecurity breaches for many healthcare organizations. Most medical professionals do not have a background in information technology or cybersecurity, yet they are responsible for assessing which treatment provides the best balance of risk and probability for success. This paper presents a two-part risk assessment framework that uses a doctor's worst case assessment of a device's potential to impact a patient and a security questionnaire based on the STRIDE model to generate a risk score on a scale from 0 to 10. Four test cases based on relevant medical devices are used to demonstrate the practical application of the framework.

15. SUBJECT TERMS

Cybersecurity, Medical, Device, Risk, Scoring

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			LTC Mason Rice, AFIT/ENG
U	U	U	U	71	19b. TELEPHONE NUMBER *(include area code)* (937) 255-3636; mason.rice@afit.edu