Securing Hospitals

A research study and blueprint

Independent Security Evaluators | www.securityevaluators.com

February 23, 2016

Abstract

The research results from our assessment of 12 healthcare facilities, 2 healthcare data facilities, 2 active medical devices from one manufacturer, and 2 web applications that remote adversaries can easily deploy attacks that target and compromise patient health. We demonstrated that a variety of deadly remote attacks were possible within these facilities, of which four attack scenarios are presented in this report. To understand these ecosystems, a two year study was performed from January, 2014 through January, 2016 of critical elements within these facilities as they relate to securing patient health. Our goal was to create a blueprint –a step-by-step action plan– that all medical facilities can follow as the foundational element in reaching full security readiness. The research was driven by a hands-on analysis of various healthcare systems, applications, and budgets, interviews with hospital, data center, and medical device manufacturer employees, and sourcing industry knowledge from thought leaders on our advisory board. The findings show an industry in turmoil: lack of executive support, insufficient talent, improper implementations of technology, outdated understanding of adversaries, lack of leadership, and a misguided reliance upon compliance. These findings illustrate our greatest fear: patient health remains extremely vulnerable. The output of the research is the production of a modern patient-health focused attack model, and a blueprint that advocates a phased approach to security design and implementation for healthcare facilities that focuses on the protection of patient health assets.

Executive Summary

This report delivers the results of our research in investigating a variety of hospital and healthcare-related infrastructures and systems, identifying industry-specific pitfalls and shortcomings, and creating a blueprint for how entities in the space can improve their security posture by the most effective means. In all, we investigated 12 healthcare facilities, 2 healthcare data facilities, 2 active medical devices from one manufacturer, 2 web applications, and a multitude of other devices, applications, and systems found on these healthcare facility networks.

AN INADEQUATE THREAT MODEL

One overarching finding of our research is that the industry focuses almost exclusively on the protection of patient health records, and rarely addresses threats to or the protection of patient health from a cyber threat perspective. The background, motivating factors, nuances, and misunderstandings that perforate the healthcare industry with regard to security are discussed at length in this report. In summary, we find that different adversaries will target or pursue the compromise of patient health records, while others will target or pursue the compromise of patient health itself. These adversaries and their likely targets are summarized here.

	Patient Health		Patient Records	
Adversary	Targeted (Specific Victims)	Untargeted (Indiscriminate)	Targeted (Specific Victims)	Untargeted (Indiscriminate)
Individual / Small Group				YES
Political Groups / Hacktivists /			YES	
Organized Crime	YES		YES	YES
Terrorism / Terrorist Org.	YES	YES		
Nation States	YES	YES	YES	YES

The two major flaws in the healthcare industry with regard to threat model are that 1) the focus is almost entirely on protecting patient records, and 2) the measures taken address only unsophisticated adversaries: essentially, only one of the adversaries listed above –the Individual or Small Group adversary highlighted above in yellow. The industry is aware and speaks to Organized Crime and Nation State adversaries, but underestimates their sophistication and motivation. The strategies aim to curtail blanket, untargeted (i.e., indiscriminate) attacks to obtain patient healthcare records, and ignores the motivations and strategies that would be employed if targeting patient health or specific victims' health records. These motivations and scenarios are highlighted in red in the above table.

As a result, a multitude of attack surfaces are left unprotected, and attack strategies that could result in harm to a patient are not considered. The following summary provides an overview of these types of attacks.

PATIENT HEALTH ATTACK MODEL

One of the primary contributions of this research is the Patient Health Attack Model. To our knowledge, no comprehensive attack model is available for the healthcare industry that catalogs the attack surfaces affecting patient health assets. We have cataloged and describe in detail the following primary, secondary, and tertiary attack surfaces that expose patient health. The following diagram illustrates this attack model.

Primary Attack Surfaces

- Clinicians
- Medicine
- Active Medical Devices (AMD)
- Surgery

Secondary Attack Surfaces

- Patient Samples
- Passive Medical Devices (PMD)
- Electronic Health Records (EHR)
- Test Results
- Work Orders
- Connected Power
- Schedules
- Inventory Systems
- Sanitary Conditions
- Procedure Precision
- Time

Tertiary Attack Surfaces

- Inventory Systems
- Climate Controls
- Environmental Controls
- Physical Storage
- Physical Transport
- Barcode Scanners / Printers
- Connected Power
- Laboratory Equipment
- Clinicians

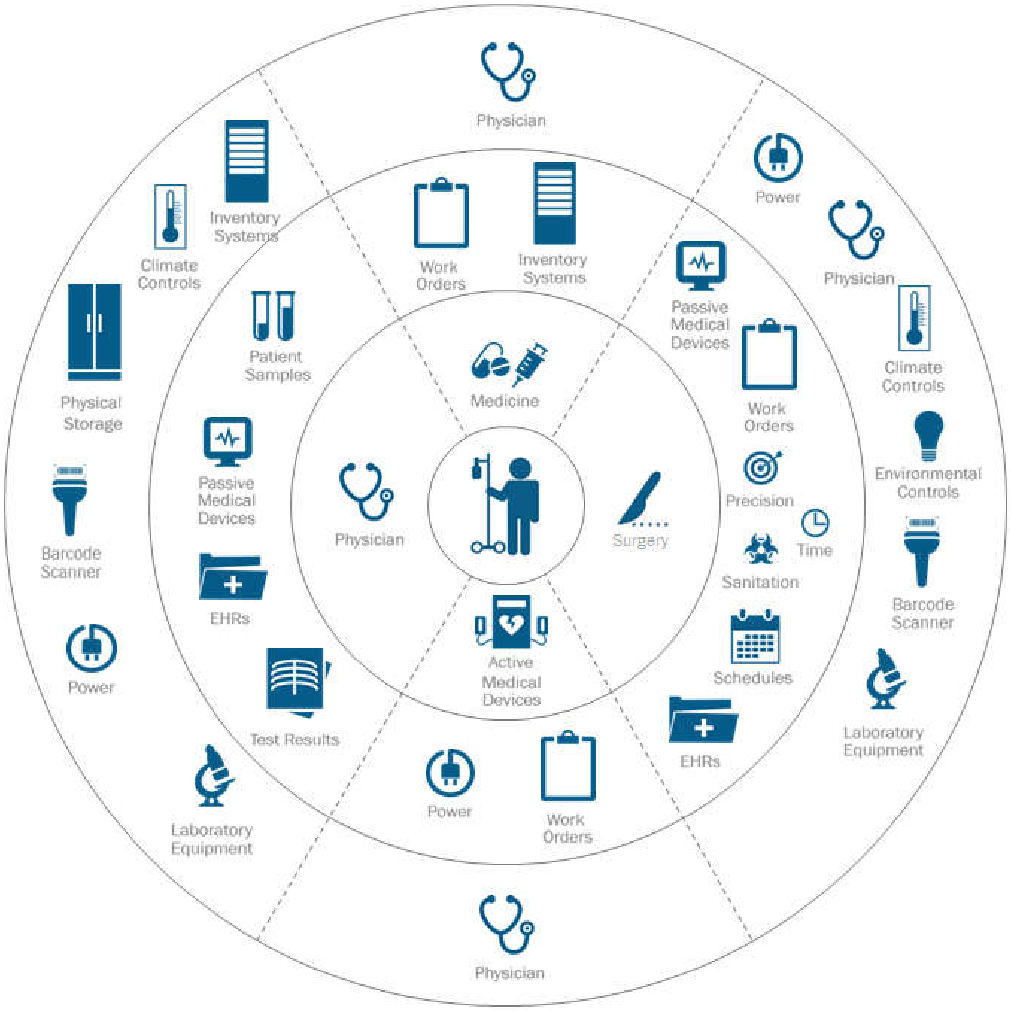

Many of the above attack surfaces have little value with regard to personally identifiable information (PII) or personal health information (PHI)–the assets hospitals strive to protect most—yet they have direct consequences with regard to patient health. These attack surfaces are largely left unprotected by hospitals and are precisely the attack surfaces to be targeted by an adversary seeking to harm a patient.

COMMON DESIGN ISSUES

We found that the hospitals were failing on a variety of levels to properly address modern security threats. Problems ranged from business-level, organizational problems (e.g., a lack of funding, staff, or training) to technical problems specific to departments (e.g., vulnerable network design, use of legacy systems, and the use of vulnerable vendor systems). Through the fog, it is difficult to pinpoint which issues are the impetus for others, as many of the problems directly or indirectly exacerbate the others, amplifying issues. However, the issues listed in the table to the left are the most notable security design deficiencies with hospitals we investigated.

Business

- Lack of funding
- Lack of appropriate staffing
- Lack of effective training
- Improper organizational structure

Policies and Procedures

- Lack of defined policy
- Lack of audit procedures

Technical

- Lack of network awareness
- Lack of logging/monitoring
- Insecure network architecture
- Insufficient access controls
- Extensive use of legacy systems
- Inability to assess/patch
- AMDs on non-restricted subnets

Vendors

- Weak remote access controls
- Use of insecure vendor systems
- Use of insecure custom systems

Physical Security

- Guest phys. access to systems
- Guest phys. access to networks
- Credentials exposed to guests

We believe the impetus for most security issues in hospitals stems from a drastic lack of funding for security departments, a lack of appropriate staffing of security personnel, and a lack of effective security training at all levels of the organization. Until these issues are addressed, it will be difficult to overcome some of the other design flaws.

Hospitals had very few proper security policies and procedures, and those that did exist were ineffective in practice. Furthermore, very little was done with regard to audit to determine what security problems existed and to create action plans to address them. Without proper policy and procedures in place, it will likely lead to heavy waste and the implementation of ineffective technical security measures.

With regard to technical security design issues, we found that hospitals were antiquated in their network designs, and unsure about the technologies that could effectively help them. In many cases, vendor products purchased for a security purpose were inappropriate for the organization, and those systems that were appropriate were deployed incorrectly, all resulting in heavy waste while not achieving an improvement in security posture. These issues were compounded by the fact that numerous vendor-installed and in-house built systems we investigated were rife with security vulnerabilities.

Hospitals also face a variety of unique problems that require special attention when addressing. Untrusted parties (i.e., patients and visitors) often have physical access to equipment and networks. People are an asset in these facilities, which is uncommon in most organizations' security models. Furthermore, time, accuracy, and environment play a role in the survival of those assets –a circumstance not found in many other scenarios.

RECOMMENDATIONS

The resolutions for these issues are not trivial. They will involve effort and diligence at all levels within the healthcare industry. In some cases, it may take years for a single hospital to reach an appropriate level of security readiness. Likewise, it will take the industry several years to correct systemic issues and create effective programs for bolstering security on every level, from the device vendor, to the hospital, and to the patient at home.

The industry should course correct to drive change toward an overall stronger security mindset. It is the responsibility of all parties involved to participate honestly and strive for the best interests of the end users: patients. For healthcare facilities, there is no question that the ultimate priority *is* to protect patient health.

For the Industry

Focus on patient health

The industry is hyper-focused on protecting patient data, which while important should come second to protecting patient health.

Avoid (or create effective) regulations

For almost two decades, HIPAA has been ineffective at protecting patient privacy, and instead has created a system of confusion, fear, and busy work that has cost the industry billions. Punitive measures for compliance failures should not disincentivize the security process, and healthcare organizations should be rewarded for proactive security work that protects patient health and privacy.

Empower the consumer

An industry-wide comparative security ranking system would empower the consumer to make informed decisions about the security of their health and privacy when choosing a provider.

Empower the CIO/CISO and other executives

Decision makers at healthcare facilities have little insight or control over the security practices of their vendors. Third-party security assessments by experienced professionals can lend to empowering the CIO and other executives if vendors are required to produce such evidence.

Philanthropy

Good security is often cost prohibitive. Much like an endowment, grant, or donation of funds that could be used for medical equipment or staffing, these funds can be appropriated to elevate the security posture of an organization.

For Hospitals

Follow the blueprint

In this report we've included a blueprint for better healthcare facility security. This blueprint should be adopted by the organization.

Create a long-term plan

Long term security plans should be understood at the executive and board levels within the organization. They should address immediate and long term efforts, including financial, staffing, training, and technology plans. Plans should be updated and evolve over time.

Increase funding

We identify the lack of good security in healthcare facilities as being heavily influenced by a general lack of funding to these areas across organizations. Nearly all aspects of the blueprint require a budget allocation to be successful.

Increase security knowledge

The facility should endeavor to increase its overall security knowledge through training, and augmenting their team with seasoned security professionals or outside consultants who can competently design and execute a security strategy.

Separate Info. Security from Info. Technology

While both areas involve technology, it is inappropriate to treat Information Security as an Information Technology effort. Information Security should separate from Information Technology with independent reporting structures at the Board level.

SECURITY BLUEPRINT

For most healthcare facilities, it is not a question of *am I secure*, or *how secure am I*, but of *how do I get there*? This question of how to get from where they are to a point of security readiness is difficult, and the further that distance the more daunting this task becomes. When the task at hand is discouraging, it is prone to delay, waste, and failure. We provide this blueprint (summarized here) for healthcare senior executives responsible for information security and patient care.

Each of the below phases are described in detail in the last section of this report. The entire process is cyclical, but each phase builds on the output of the previous phases. Each phase and sub-step is essential, though we found that most healthcare organizations focused only on a very small subset of these stages, and often late stage exercises only; these late stage exercises proving to be of little overall effect given they were not preceded by the appropriate planning or design steps.

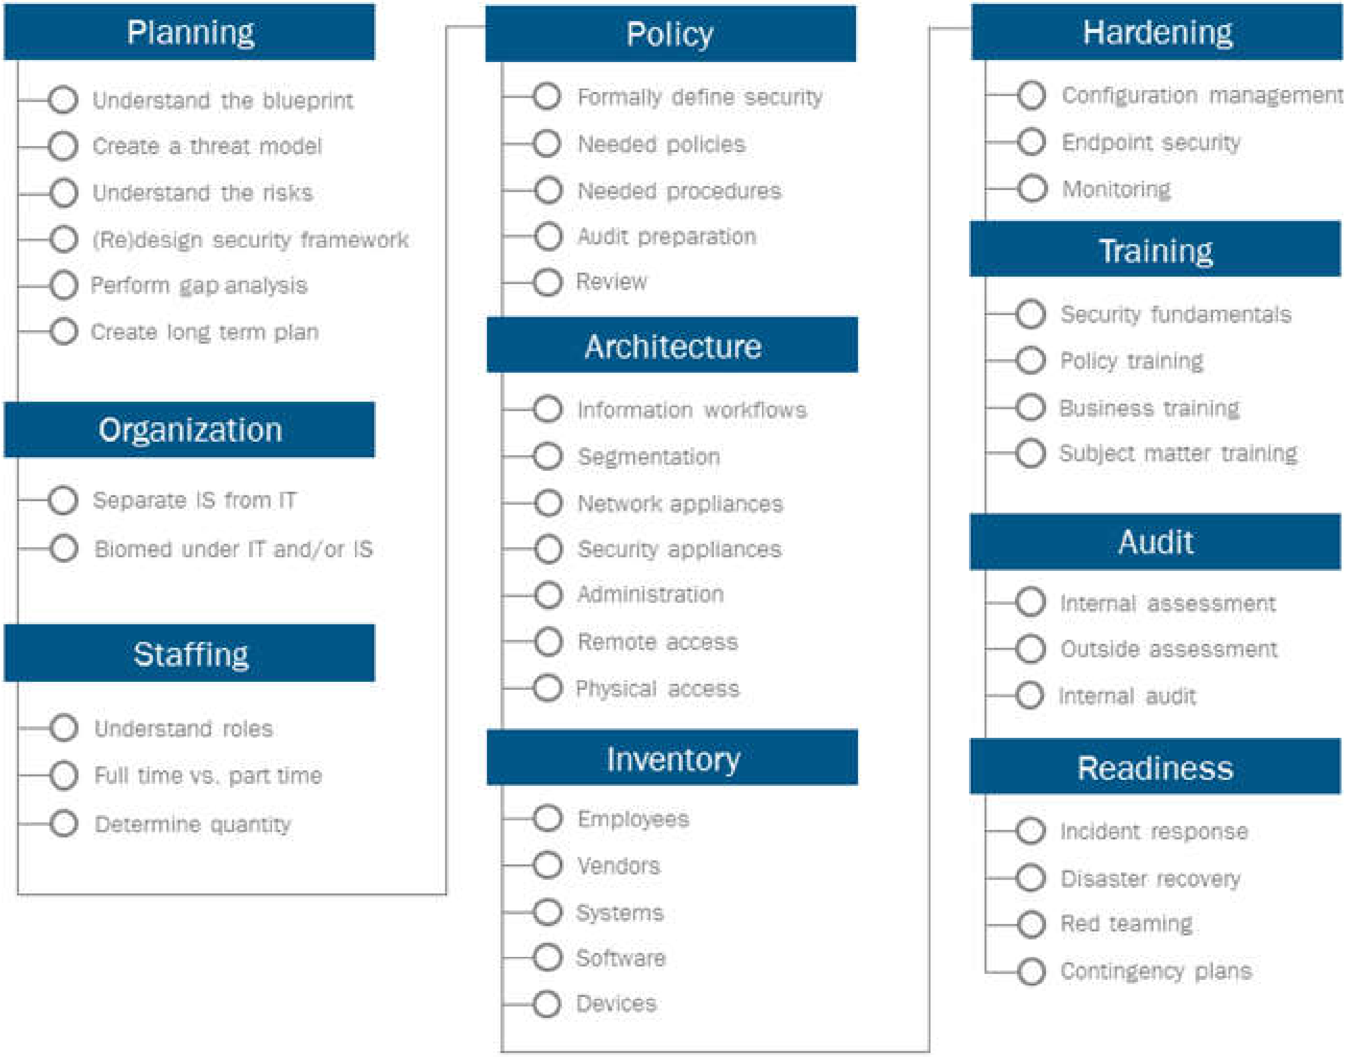

Table of Contents

Part I: Background and Introduction

We hope that this research can both raise awareness and direct future efforts toward creating a safer and more secure healthcare technology infrastructure. To date, we know of no real-world attacks against individuals or groups of patients, but our findings discussed throughout this report suggest that these attacks are readily possible and have the propensity to succeed in causing physical harm to patients in most healthcare settings.

We believe these attacks against patient health are real and present, and likely to be acted upon in the near future. Research in the security community has demonstrated repeatedly that medical devices can be compromised and controlled to cause harm to those patients to which they are connected. As evidenced by extensive news reports and our own observations of the medical field that are further confirmed by our research here, it has been demonstrated that the infrastructures surrounding these devices are vulnerable. This represents opportunity. Motive is beyond the scope of our research, but we lean on the de facto assumption that organized crime, terrorism, and nation state enemies have the motivation to cause physical harm to patients enrolled in the healthcare systems of the entire world. With both motive and opportunity, we anticipate attacks will be realized and highly disruptive.

We are motivated in this research because these threats to patient health are threats to our individual selves, our families, our communities, our economy, and our national security. We hope that this research and our suggestions are adopted industry-wide in efforts to create a secure healthcare industry.

Heading in the *wrong* direction

The mission of security in healthcare is focused on protecting patient health records, and ignores patient health. This is evidenced openly in legislature through HIPAA, HITECH, and other legislation and regulatory directives that command fines in response to the loss of patient records, but speaks sparingly to patient health. As a result, this drives internal directives to focus on protecting these records, but offers little guidance or incentive for protecting patient health. The efforts that do aim to protect patient health do not address intelligent cyber threats. Defending patient health and patient records is not one-in-the-same, and placing the focus on records harshly ignores the patient health aspect. So long as this is the mission of the industry, it is unlikely that patients' health will be adequately protected in the healthcare ecosystem.

Wrong mission **x** **Outdated approach** **=** **Failure**

Focusing on patient records *Ignoring advanced threats* *Patients not protected*

Furthermore, the mission to address even the records aspect of these issues considers an outdated and inappropriate adversary. The driving efforts focus almost entirely on unsophisticated, untargeted attack areas, such as wide-scale data loss prevention –a truly important initiative, but incomplete when faced by legitimate, sophisticated adversaries. Such sophisticated attacks are very real and evidenced in other industries. To simply focus on the lowest bar of protection does a disservice to patients who remain unnecessarily exposed to those adversaries willing to put forth a slightly greater effort. This is a common fallacy that has been realized and addressed in other industries, and must now be addressed in the healthcare space.

Regulation across many industries, including healthcare, has sought to reduce the threat from adversarial compromise, but they have only been successful at reducing the damages from those adversaries in the least sophisticated, untargeted categories. We believe that healthcare relying heavily on regulation as the saving motivation for protecting patient records or health is also seriously misguided, and will not result in a safer or more secure healthcare ecosystem for patients' health, privacy, or identity.

Challenges to success

One can easily observe the disarray and indicators of unlikely success, heavy waste, and poorly directed efforts. There is blanket criticism of regulatory statutes among security professionals, and statistics have been showing dramatically increasing losses, not successes. Digging further, it is evident the causes of these increased losses. Hospitals have severely marginalized budgets with very little focus on security. Perhaps as a result of this, we routinely encounter undertrained and understaffed teams; often with hospital security teams having **zero** information security personnel. Until this process is course-corrected, losses and waste will increase.

There are significant challenges in changing trajectory. First, capable security talent is hard to obtain. The demand for information security professionals far outweighs the supply and there is arguably a 0% or negative unemployment rate in this sector. Experienced, seasoned talent is even harder to obtain; and then, no one is left to make the determination of talent fitness. Until appropriate security professionals exist within an organization, it will be very difficult to secure that infrastructure or for the decision makers to understand the threats they face. Second, the healthcare information technology market is perforated by misunderstood and misrepresented service and product offerings. Term confusion and the promises of pipe dream (turn-key) solutions foster waste and false confidence. The healthcare community is in need of legitimate, actionable steps that can be followed to obtain stronger and more secure security postures.

Hospital Challenges

Lack of budget

Understaffed

Undertrained

Heavy waste

Industry Challenges

Regulatory interference

Misrepresented services

Lack of talent

Lack of direction

A solution

Our goal is to provide an effective and actionable blueprint for correcting this trajectory on a case-by-case basis. Hospitals have unique problems that are not applicable to traditional business, and thus require unique solutions. Patient health assets exist in very few other industries and regulation is stringent in healthcare unlike many other industries. It is not reasonable to simply adopt the methodologies of other industries and apply them to healthcare. Within healthcare, however, hospitals certainly face the same regulatory, budgetary, organizational, political, public perception, and day-to-day work flow issues. This justifies a uniform blueprint approach.

While a blueprint is not an end-all solution to security in any industry, they have a number of benefits. They provide a solid foundational security plan, and allow less experienced, less trained individuals to benefit from the findings of more experienced, seasoned security professionals for whom they may not have access. A blueprint can prevent adopting less effective means, reducing both waste and delay, and can help justify budgets and quantify risk-reward estimations, reducing both waste and risk.

This research provides a blueprint as a starting point, and not a turn-key or end-all solution to the security problems faced by healthcare. Hospitals and other healthcare organizations who cannot obtain the requisite security personnel should continue to seek outside expertise to help harden their infrastructure and create long term security plans and audit against them.

Introduction

This report delivers the results of our research in investigating a variety of hospital and healthcare-related infrastructures and systems, identifying industry-specific pitfalls and shortcomings, and creating a blueprint for how entities in the space can improve their security posture by the most effective means.

First, we provide a background of participants involved in this research. Next, we describe our methodology and provide a modern threat model by which our research was conducted –and by which all patient-focused security programs should be designed. We describe some of the real-world attack scenarios we uncovered. We discuss general design issues with hospital infrastructure security, and recommend solutions. Lastly, a blueprint is provided by which healthcare organizations can benefit as a starting point to becoming more secure.

This report is not a comprehensive survey of the industry, nor does it represent a one-size-fits-all solution to security should the blueprint be followed. It is meant to be a starting point, and justification for a change in the trajectory of the industry. It is important to continuously recognize that even with a proper plan in place, proper execution of that plan is essential in order to reach the goal: a more secure infrastructure that addressing securing patient health. This research provides the scaffolding for that plan.

The blueprint portion of this report can be adopted by hospitals to begin planning for security infrastructure revisions. The security team, in concert with the executive decision making bodies of these organizations should review the blueprint and decide on which aspects are most pertinent to the organization. Those organizations who do not have sufficient expertise should seek it out.

About ISE

ISE was born in 2005 out of the PhD program at the Johns Hopkins Information Security Institute, and for over 10 years has helped enterprises protect digital assets from sophisticated adversaries by employing the same methodology and mindset perpetrated by those adversaries. ISE analysts are domain experts in the crucial security disciplines, including cryptography, reverse engineering, malware analysis, design verification, social engineering, and many more. ISE analysts bring a diversity of experience with analysts coming out of PhD and other academic programs, and others bringing industry background from esteemed security organizations across various industries.

Research team: Stephen Bono, Thomas Connolly, Paul Dant, Geoff Gentry, Ted Harrington, Jacob Holcomb, Jacob Thompson, and Larry Trotter.

Advisory Board

In conducting this research, ISE formed an advisory board of experts involved in various aspects of the healthcare field. We relied on this advisory board for expert advice and guidance during this project. The advisory board is staffed by a representative cross section of the healthcare industry, drawing upon their expertise to ensure this research could be most effectively put to practice. The board includes physicians and nurses – for medical opinion on how attacks could affect patients; lawyers – for how our suggestions exist within the scope of existing compliance and regulatory statutes; and hospital CIOs – for explanation of hospital day-to-day operations and set-backs.

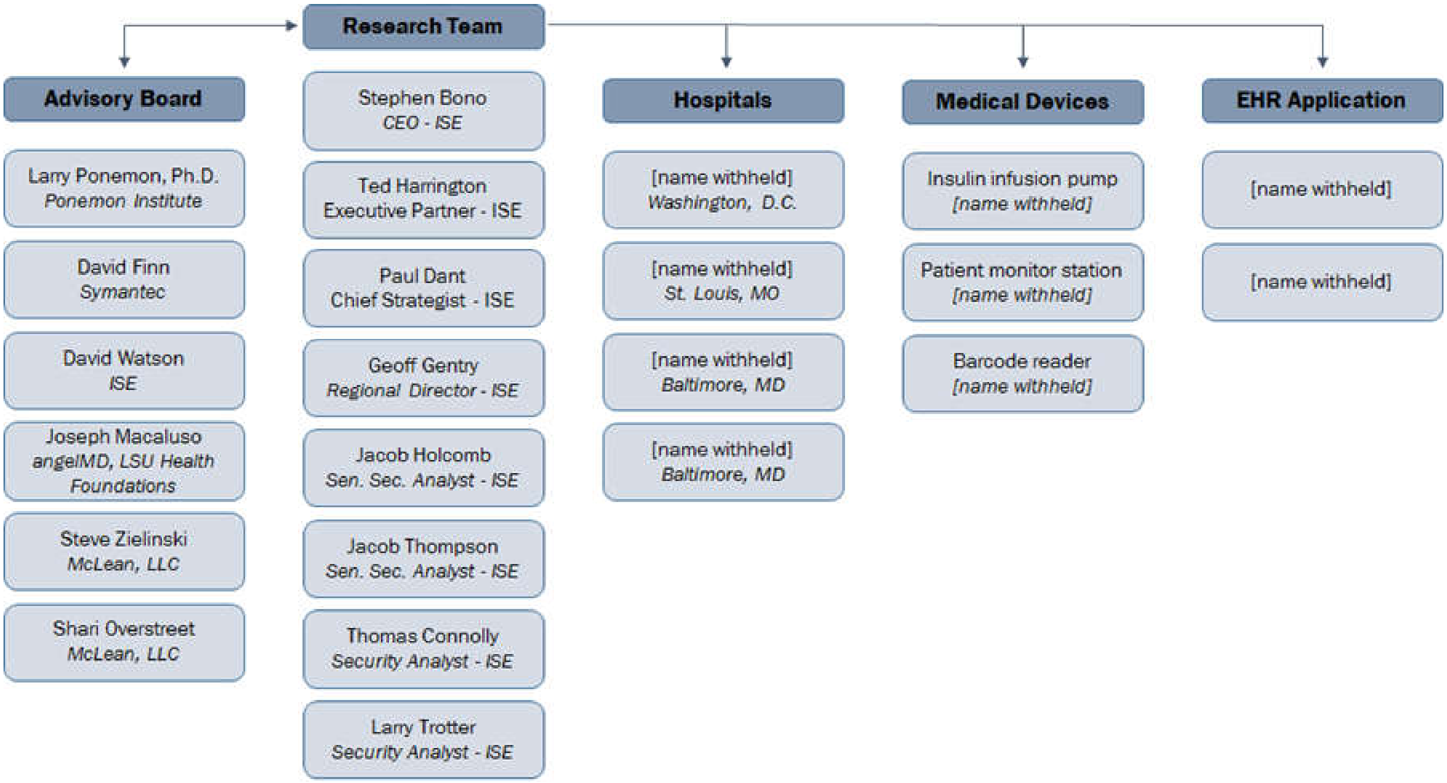
Research Team
Advisory Board
Larry Ponemon, Ph.D. Ponemon Institute
David Finn Symantec
David Watson ISE
Joseph Macaluso angelMD, LSU Health Foundations
Steve Zielinski McLean, LLC
Shari Overstreet McLean, LLC
Stephen Bono CEO - ISE
Ted Harrington Executive Partner - ISE
Paul Dant Chief Strategist - ISE
Geoff Gentry Regional Director - ISE
Jacob Holcomb Sen. Sec. Analyst - ISE
Jacob Thompson Sen. Sec. Analyst - ISE
Thomas Connolly Security Analyst - ISE
Larry Trotter Security Analyst - ISE
Hospitals
[name withheld] Washington, D.C.
[name withheld] St. Louis, MO
[name withheld] Baltimore, MD
[name withheld] Baltimore, MD
Medical Devices
Insulin infusion pump [name withheld]
Patient monitor station [name withheld]
Barcode reader [name withheld]
EHR Application
[name withheld]
[name withheld]

LARRY PONEMON, PH.D. – PONEMON INSTITUTE

Dr. Larry Ponemon is the Chairman and Founder of the Ponemon Institute, a research “think tank” dedicated to advancing privacy and data protection practices. Dr. Ponemon is considered a pioneer in privacy auditing and the Responsible Information Management or RIM framework.

Ponemon Institute conducts independent research, educates leaders from the private and public sectors and verifies the privacy and data protection practices of organizations in a various industries. In addition to Institute activities, Dr. Ponemon is an adjunct professor for ethics and privacy at Carnegie Mellon University’s CIO Institute. He is a Fellow of the Center for Government Innovation of the Unisys Corporation.

DAVID FINN – SYMANTEC

David Finn, CISA, CISM, CRISC is the Health Information Technology Officer for Symantec. Prior to that role he was the Chief Information Officer and Vice President of Information Services for Texas Children’s Hospital, one of the largest pediatric integrated delivery systems in the United States. He also served as the Privacy and Security Officer for Texas Children’s. Prior to that Mr. Finn spent 7 years as a healthcare consultant with IMG/Healthlink and PwC, serving last as the EVP of Operations for Healthlink.

Mr. Finn has more than 30 years’ experience in the planning, management and control of information technology and business processes. He is focused on enabling operating efficiency and deriving business value through the optimization and control of technology. Mr. Finn’s key skills include IT Governance and Control, Project Management, Systems Selection and Implementation, Business and IT Partnering, and IT Audit, Control and Security.

In addition to having served on the national Board of HIMSS, he currently serves on the CHIME Board of Trustees. During 2014, Mr. Finn worked closely with CHIME management to create and initiate the Association for Executives in Healthcare Information Security (AEHIS). In the past, he served on the Information Systems Audit and Control Association's (ISACA's) Professional Influence and Advocacy Committee (PIAC). He also is a long-time Board member of Healthcare for the Homeless - - Houston (2 FQHCs) and is Vice President of the Primary Care Innovation Center in Houston.

DAVID WATSON – INDEPENDENT SECURITY EVALUATORS

David Watson has a vast array of experience in network infrastructure management, architecture and design, data management, application management and is a security program manager for Independent Security Evaluators (ISE.) Previously he was a portfolio manager for Ascension Health, the nation's largest non-profit healthcare system. During this time he was responsible for overseeing the largest private health information exchange (HIE) in the state of Texas, as well as, the business intelligence and analytics program for Seton Family of Hospitals based in Austin, TX. Prior to joining Ascension Health, David was an independent consultant focused on healthcare information technology program management. David has sat on advisory boards for University of Texas' new data center build, HIMSS Enterprise HIE task force, and has been a Director for Young Professionals in Energy.

JOSEPH B MACALUSO, JR., M.D. FACS - ANGELMD, LSU HEALTH FOUNDATIONS

Dr. Macaluso has a long history of accomplishment in medicine, clinical practice, surgery and urology. He maintained one of the most active urological surgery practices in the nation for more than 22 years and served as the Managing Director and Director of Research and Grants at the Urologic Institute of New Orleans for 15 years. He taught medical students and residents for many years as an Associate Professor of Clinical Urology at Louisiana State University Medical School and Charity Hospital in New Orleans. Dr. Macaluso also held the rank of Assistant Professor of Clinical Urology at Tulane Medical School. Board certified by the American Board of Urology, Dr. Macaluso has been cited by numerous "Best Doctors" lists, and was named repeatedly in Best Doctors in America. His dedication and commitment to quality patient care and research is well known throughout the urology profession. He retired from active practice in 2005.

STEVE ZIELINSKI – THE MCLEAN GROUP

As a leader of financial services firms, Mr. Zielinski is the Managing Director at The McLean Group's St. Louis, MO office. His 25 years' experience as a financial professional, management consultant and investment advisor to middle market institutions and businesses have provided him a strong background in buy-side and sell-side investment banking transactions.

As a former president and chief investment officer of a financial services firm, Mr. Zielinski has focused on using innovative tools and approaches to obtain financing from institutional and accredited investors, and government sources to fund ventures in cleantech, biotech, healthcare, education and agribusiness.

SHARI OVERSTREET – MCLEAN L.L.C

Ms. Overstreet has been a Finance and Accounting professional for over 30 years. She holds a CPA license and is a FINRA licensed investment banker. She also holds a variety of business valuation, and merger and acquisitions-related designations. During her career, Ms. Overstreet has worked for large accounting firm, Arthur Andersen, as an auditor and tax professional. She has also served in positions such as controller, Director of Finance, and Chief Financial Officer for companies both publicly-traded and privately-held, whose annual revenues ranged anywhere from $1 Billion per year to smaller, start-up companies.

Ms. Overstreet is an author and speaker on various mergers and acquisitions, capital formation and business valuation topics. She was a 2010 Nominee for the Profiles in Power & Women of Influence of Central Texas Award. She holds a BBA with a finance concentration from the University of Texas at Austin.

Participants

Our research targeted medical facilities in the following locations:

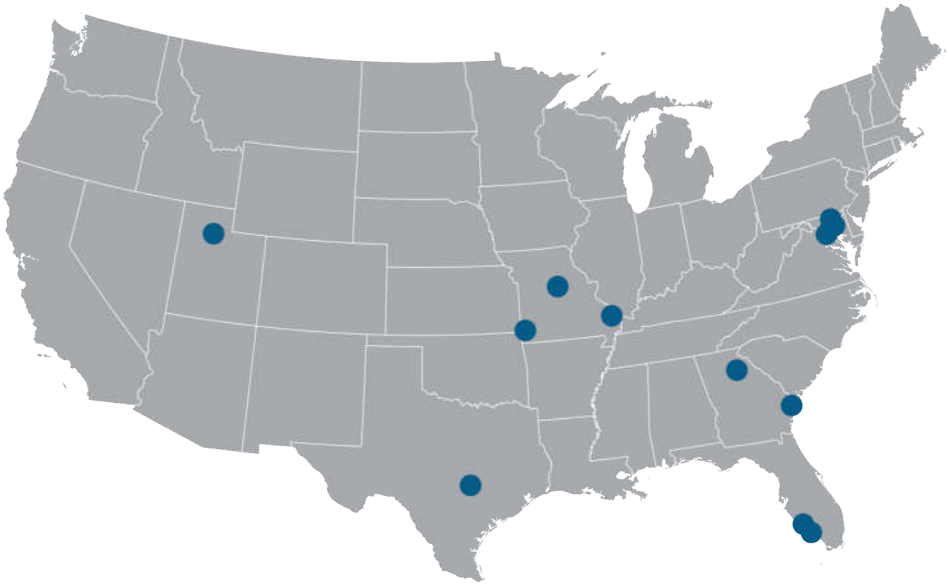

Baltimore, MD
Towson, MD
Washington, D.C.
Athens, GA
Savannah, GA
Cape Girardeau, MO
Columbia, MO
Joplin, MO
Salt Lake City, UT
Naples, FL
Bonita Springs, FL
Austin, TX

Additionally, our research targeted a multitude of devices and applications, including a variety of in-house developed and commercial Electronic Health Records management systems.

As we progressed through our research, we investigated numerous components that were originally out of scope. Many of these components provided valuable intelligence to the overall efforts in performing our research, and are woven throughout our findings, although perhaps not mentioned specifically.

Threat Model

Effective risk management requires an understanding of both the system to be defended and the adversaries that threaten it. Assets that require protection need to be identified along with the impact a successful attack would have on those assets. Threat actors' intentions and capabilities need to be modeled and applied against vulnerabilities, and their likelihood to impact the identified assets. This allows for informed decisions about which available mitigations to apply resources against, and results in the most secure systems possible within potential resource constraints.

Assets

The following are the primary assets found within the healthcare ecosystem. First listed are patient-specific assets. Patient health, in particular, is listed at the head of this category and should be considered the highest priority asset to protect. Other assets may indirectly affect patient health. Second listed are hospital and other organization-specific assets. These do not affect patients in any way as directly as the patient assets, but play an indirect role. Attacks against hospital assets can 1) indirectly disrupt patient care, 2) raise the cost of healthcare, and 3) hinder the progression of the industry toward beneficial care potential.

Patient Assets	Hospital Assets
Patient health	Research / IP
Patient records	Buisness advantage
Service availability	Hospital finances
Community confidence	Hospital reputation
	Physician reputation

PATIENT HEALTH

Patient health must be the paramount asset of greatest importance to protect within the healthcare industry. "First do no harm" is a motto adopted by healthcare professionals, and this should be extended to the practice of security by those supporting them. Patient health could be affected in a variety of ways, including causing permanent or temporary physical or mental injury, disrupting care in some way so that treatment cannot be obtained, and even causing death.

PATIENT RECORDS

Patient records are incredibly valuable to patients and adversaries alike. They include private information that the patient and others may not desire to be made public, and they are of high value to identity thieves who may wish to abuse the information contained within for financial gain. Patient records may include personally identifiable information, such as social security numbers, health care provider information, credit card information, name, address, date of birth, etc. Records may also include the private health information about a patient's mental or physical health or the patient's social history. Records also play in to patient health through integrity; if records can be altered or destroyed, it could adversely affect patient health.

SERVICE AVAILABILITY

Attacks on healthcare service availability can be devastating to both patients and providers. These attacks could prevent critical services which can lead to patient injury, but also to deny service for the purposes of paying bills, filling prescriptions, making appointments, or getting help. There is relatively little to gain for the adversary in doing so, but nevertheless these attacks do occur and can be serious.

COMMUNITY CONFIDENCE AND TRUST

Should patients or the community lose trust or confidence in the healthcare industry's ability to help them, or become afraid to engage them, it could undermine the overall health and safety of our country. Examples of widespread loss of confidence that have had negative effects on our safety and economy can be seen in the distrust of airport security following the attacks of September 11, 2010, or communities developing distrust for police or government after specific incidents arise or appear to arise. Similar phenomena have also occurred in healthcare, as can be seen in the sudden wide-spread distrust by parents of child vaccinations. If similar widespread loss of confidence were to afflict the healthcare industry, such as the community refusing to seek treatment due to fear of harm (justified or not), it would be extremely detrimental to our health, safety, and economy.

RESEARCH AND DEVELOPMENT / INTELLECTUAL PROPERTY

Of less concern to patients, but very real within the healthcare ecosystem are the intellectual property assets that make up research and development efforts at hospitals. These could be drug formulas, test results, surveys, test subject information, experimental procedures for surgery, large scale analytics databases, etc., all of which represent high value to owner and adversary alike. Unless involved in a drug trial of some kind, patients are unlikely to be concerned with this asset.

BUSINESS ADVANTAGE

Hospitals are not just healthcare providers, but are also businesses with competitors, strategies, market share, and some are even publicly traded on the stock market. This provides a high value opportunity for corporate espionage and other malicious actions that could give one hospital or organization advantage over another. These assets are valuable to both adversary and hospital alike, and are likely to be the target in cyber-attacks today and in the future.

HOSPITAL FINANCES

Much as the theft of personally identifiable information (PII) to an adversary has significant value on the black market, so do attacks against the hospital as a financial entity as well. Like any business, the hospital may be targeted to obtain employee payroll records, corporate bank account records, or accounts payable and receivable information in order to abuse them for financial gain.

HOSPITAL REPUTATION

Hospitals and healthcare providers place enormous value in their brand and reputation. It would be a serious oversight to ignore the fact that protecting patients, their records, and their partners' research and development efforts have a direct correlation to the providers' reputation should those assets be compromised.

PHYSICIAN REPUTATION

Physicians, like hospitals, are scrutinized and weigh heavily their reputation for success. Attacks that could intentionally or indirectly affect a physician's reputation, such as impersonating a physician in an attack, compromising a physician's workstation, or leveraging a physician's stolen credentials in an attack could all result in that physician losing credibility or suffering harm to their reputation.

Understanding Adversaries

Before diving in to adversary specifics, it is important to address the following concerns; now confirmed by our research:

1) The failing of healthcare facilities to account for both untargeted *and* targeted attacks, and

2) The failing of healthcare facilities to account for both unsophisticated *and* advanced attacks.

The action taken by the industry thus far is largely reactionary focusing on addressing the many unsophisticated, untargeted attacks that have plagued the healthcare industry. By ignoring the motivation for and evolution of these attacks and focusing only on the symptoms, it has furthered a security approach that–even if ever successful against the present threats—will fail as threats evolve to the next level.

CRIME AS A BUSINESS

Attacks on healthcare are prevalent not simply because the attacks are easy; instead, attacks are prevalent because the assets available for compromise have high value to those adversaries performing the attacks. Cybercrime is a lucrative business, and as long as the costs of performing an attack are less than the expected gains, the attacks will continue. Indeed, the most likely attacks will come when the difference between cost and reward are greatest, but this nuance of adversarial motivation is often overlooked. As a result, security focus in healthcare is applied to the symptoms –the specific nature of the latest, known breach– and a fantastic false sense of security arises from the perceived downturn in attack activity.

One must recognize that adversaries are motivated by gain. This dictates behavior and is a predictor of the future. Consider the following, simple condition weighed by the adversary before an attack is launched:

(1) If *cost* + *risk* < *reward,* ***do it***, else ***don't do it.***

This is a simple business value proposition. Where the healthcare field has failed in constructing adequate security measures comes from two corollaries to this condition:

(2) Given two attacks, if the *cost* and *reward* are the same, ***choose the lower risk attack.***

(3) Given two attacks, if the *risk* and *reward* are the same*,* ***choose the lower cost attack.***

These rules tell us that attacks will not necessarily stop as defenses improve, but instead evolve. The security posture of most healthcare facilities is not prepared for an evolving adversary.

UNTARGETED VS. TARGETED ATTACKS

Whether an attack is targeted or not depends on the adversary's motivation. Untargeted attacks do not discriminate between assets, while targeted attacks have specific assets in the crosshairs. A patient electronic health record (an EHR), particularly the personally identifiable information (PII) found within that record that can be used for the purposes of identity theft and other insurance fraud opportunities, is generally not worth distinguishing between other assets of its kind. The average EHR is valued on the black market at over $50 per record[1]. To the adversary interested in selling or leveraging mass quantities of PII found in EHRs, the adversary seeks to compromise the records of *any* patient because the records have relatively equal value. This makes the attacks **untargeted**. Less common is an adversary targeting the EHR of a specific individual or group of individuals. This situation seeks to exploit the personal health information (PHI) details of the record, possibly to extort or embarrass those targeted. The value in doing so could be much greater on a per-record basis. The point being, it is readily apparent within the healthcare industry that the motivations for these attacks are vastly different.

Wrong Approach, No. 1

By focusing solely on defending against untargeted attacks, attacks against patient health are ignored. This is the current approach within the industry, and it is inappropriate when defending patient health assets.

Untargeted attacks have advantages. The lack of discrimination means that adversaries can choose the weakest targets first. This could mean the weakest infrastructure (targeting one insecure hospital over its more secure neighbor), or the weakest attack surface (targeting a hospital's externally facing EHR portal over a multi-phase attack campaign to compromise an internal database). Additionally, untargeted attacks can benefit from opportunistic exposures. A lost mobile device, a password disclosed in an entirely separate breach, or simply stumbling upon EHR unwittingly can lead to the exposure of thousands of EHR with relatively little difficulty. These types of exposures aiding in the compromise of a specific, targeted asset are not likely, and thus targeted attacks are more difficult to carry out successfully –but they are possible.

Defending against targeted vs. untargeted attacks should be approached differently. There is certainly overlap in the techniques, but it is inappropriate to believe that addressing one inherently addresses the other; it does not. As the industry pursues a security approach that only addresses the untargeted adversary's motivations, it will leave open the opportunities for targeted attacks. Since a targeted attack is the most likely scenario when patient health assets are considered, this is problematic to the mission of protecting those assets.

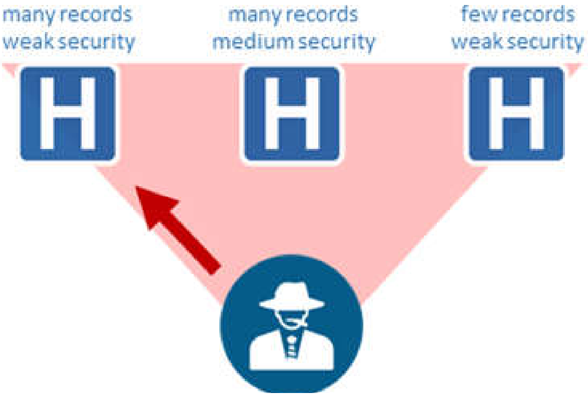

Untargeted: Adversary chooses hospital based on highest *reward* to *cost* ratio.

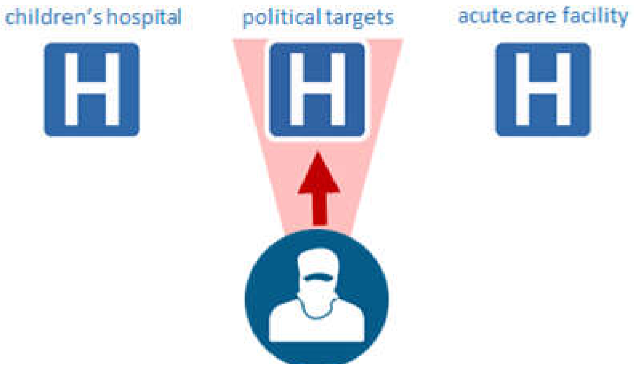

Targeted: Adversary chooses hospital based on desired victims.

[1] http://www.medscape.com/viewarticle/824192

UNSOPHISITICATED VS. ADVANCED ATTACKS

There are certainly many qualities to an attack that could make it considered either unsophisticated or advanced, but for the sake of this paper we make two important distinctions. The first is that unsophisticated attacks leverage known vulnerabilities —that is, vulnerabilities that have been previously disclosed in the afflicted systems— or are easily detected using automated tools. Advanced attacks are those that leverage 0-day vulnerabilities in applications. These may be vulnerabilities in systems supplied by vendors, or vulnerabilities in custom-built applications that are not easily detected by automated means. The second distinction regards how many vulnerabilities are exploited in series or as part of a longer-term campaign leading to the compromise of an asset. Unsophisticated attacks generally have one, maybe two vulnerabilities chained before reaching the goal, while advanced attacks may involve numerous 0-day vulnerabilities exploited over a long period of time before compromising one or many assets.

Wrong Approach, No. 2

By focusing solely on defending against unsophisticated attacks it does not address targeted attacks or the future of untargeted attacks, both of which will have advanced characteristics and remain unaddressed if the focus does not change.

ATTACKS	
Unsophisticated	**Advanced**
Leverages known issues	Leverage 0-day vulnerabilities
Chain ≤ 2 exploits in series	Chain 2+ exploits in series
Short term campaign	Long term campaign

Unsophisticated attacks should not be confused with unsophisticated adversaries. It is common for advanced attackers to employ unsophisticated attacks. Again, this reduces to the ease and cost of launching an attack –if unsophisticated methods prevail, there is no need for advanced techniques.

Defending against unsophisticated vs. advanced attacks is approached differently. As with untargeted vs. targeted attacks, there is overlap between the methods, but it is inappropriate to approach security believing one inherently addresses the other. In the same way, as the industry addresses unsophisticated attacks (i.e., addresses the symptoms) the opportunities for targeted attacks are left open. Since a targeted attack is the most likely scenario when patient health assets are considered, this is problematic.

A CHANGING THREAT LANDSCAPE

Traditional information security accounts for three attack surfaces: the physical, the human, and the digital perimeter. These attack surfaces are protected by three traditional means: physical security, training, and digital perimeter defenses such as firewalls and intrusion detection systems. Modern attacks, however, do not adhere to traditional attack patterns, and thus defending against them with an outdated approach is ineffective. The healthcare industry in particular succumbs to the belief that traditional security measures are sufficient. This is evident in regulatory statutes, proposals and presentations made by the security community, and our own experience in this research project and in other engagements. In the past, relying solely on these methods was not necessarily *correct*, but arguably effective given the environment at the time. Much has changed contributing to the current state, and the increasing likelihood that advanced attacks will be witnessed in the coming years.

What has changed?

Over the years, the number of viable attack surfaces has increased significantly with the prolific adoption desktop systems, laptops that leave and reenter the perimeter, mobile devices, vendor applications and other network-connected vendor devices –each step adding the exploitable attack surfaces as each circumvent the perimeter. Workflow in healthcare has also changed, warranting the inclusion of remote physician, vendor, and even patient access –each an opportunity to bypass the perimeter.

The accessibility to EHRs in general has increased dramatically over the past decade. Now, records are widely digitized with redundant availability, and patients and physicians alike insist on the collaboration and sharing of data to better serve healthcare needs. Coupled with the increased value of these assets on the black market, there is no surprise that attack persistence has increased. Crime as a business dictates that this increase in access and value will result in such attacks.

While not specific to healthcare, the general nature of modern attacks has evolved to disregard traditional perimeter security entirely. Advanced attacks often take months, and involve the compromise of numerous internal devices and the maneuvering throughout a network before reaching the desired assets of value. Furthermore, there are more highly trained bad guys today than ever before, let alone compared to ten years prior. As more and better advanced threats set their focus on healthcare, invariably the assets will be harder to defend.

Industry changes

- ↑ attack surfaces
- ↑ remote access
- ↑ digitization of records
- ↑ accessibility to EHR
- ↑ collaboration

Black market changes

- ↑ value of EHR
- ↑ number adversaries
- ↑ skill of adversaries
- ↑ technique of adversaries

What is still changing?

With regard to both patient health records and patient health, the same trends will continue. Increased attack surfaces will continue to lower attack cost. Increased asset value and availability will drive up attack reward. Increased adversarial skill will continue to lower attack risk. All of the above results in a greater disparity in the *cost + risk* and *reward* condition, meaning attacks will be more and more likely.

As unsophisticated and untargeted attacks are addressed, even if successfully, it will not fundamentally change the fact that the *cost + risk* of launching an attack is far outweighed by *reward*. It will only move adversaries who are already skilled in modern attack campaigns toward using them strategically and with greater precision against healthcare. Thus, it is a disservice to focus only on the unsophisticated and untargeted attacks, and those attacks that focus solely on patient health records, as those metrics will be overshadowed by already available, modern attack methods.

Actual Adversaries

In addition to the identification of assets, it is necessary for a healthcare organization to identify the adversaries for which they want to defend. Not all healthcare facilities are concerned with the same adversaries. For instance, a small healthcare facility in an unpopulated area may not be concerned with nation state or terrorist threats, while a metropolitan area hospital could be. Likewise, certain facilities may care for VIPs, or associate with a politicized cause, and therefore have a heightened threat from paparazzi or politically motivated threats.

By understanding the pertinent adversaries a facility can direct efforts in ways that:

1) Focus on the highest value activities that support the primary mission, and

2) Eliminate waste associated with defending against threats that are not present.

The following section describes the most likely adversaries faced by participants in the healthcare industry. For each adversary we discuss their motivation and sophistication, but call out in particular their relationship with the two primary assets discussed in this research: patient health and patient records. Different adversaries will approach the compromise of these assets in different ways, hence how they are protected will vary by adversary.

Adversaries

- individual/small group
- political/paparazzi
- organized crime
- terrorism
- nation state

INDIVIDUAL/SMALL GROUP

Individual and small group adversaries are motivated primarily by profit and notoriety. These adversaries generally rely on unsophisticated means and targets of opportunity.

Patient Health

⊘ unlikely to target

Patient EHR

untargeted; may not discriminate

POLITICAL GROUPS/PAPARAZZI

These adversaries are motivated by political gain, hacktivism, publicity, and financial gain. Objectives may be to obtain the medical records of high profile individuals for the sake of embarrassing or discrediting them, blackmail, or for sale to tabloid trade organizations. Objectives could also be to obtain the records from a specific, politically charged healthcare organization, such as attacks against the Planned Parenthood organization[2]. These adversaries have unqualified skill, and may seek out other skilled organizations to perform attacks for them. Notable attacks by these adversaries in other industries include attacks against the Obama and Romney campaigns in 2012[3], attacks to obtain personal photos of celebrities from Apple's iCloud in 2014[4], and attacks against the United Nations Framework Convention on Climate Change in late 2015[5].

Patient Health

⊘ unlikely to target

Patient EHR

⌖ targeted; may choose specific victims

[2] http://www.huffingtonpost.com/entry/hackers-launch-second-cyber-attack-on-planned-parenthood_us_55b9e270e4b0b8499b185c53
[3] http://swampland.time.com/2013/05/07/obama-romney-campaigns-subject-to-repeated-hacking-attempts-in-2012/
[4] http://www.businessinsider.com/apple-statement-on-celebrity-hacking-2014-9
[5] https://www.hackread.com/anonymous-hacks-un-climate-change-website/

ORGANIZED CRIME

These adversaries are motivated by financial gain and other related systemic criminal activities, such as extortion, blackmail, or coercion. Objectives may be to obtain the medical records of target individuals, or cause or threaten physical harm to target individuals, or simply to profit from the exploitation of untargeted EHR in volume. These adversaries are highly skilled, and have been involved in the black market trade and cybercrime business for decades. Unsophisticated organized crime groups can also solicit the force of skilled organizations. Notable attacks by these adversaries in other industries are the theft of $45 million from ATMs around the world in 2013[6], and cyberattacks against Target[7], Home Depot[8], and JPMorgan Chase[9].

Patient Health	Patient EHR
targeted; may choose specific victims	untargeted; may not discriminate
	targeted; may choose specific victims

TERRORISM/TERRORIST ORGANIZATION

These adversaries are motivated to inspire fear and cause harm—objectives that traditional information security may be unaccustomed to defending against. Objectives may be to harm or threaten the harm of one or a group of individuals. These adversaries do not typically demonstrate as high skill as organized crime or nation state actors, but as the opportunity for spreading fear is presented, these organizations may develop or leverage this skill, or seek to solicit the force from non-terrorist organizations. Notable attacks by these adversaries in other industries have been launched by ISIS[10,11] and the Syrian Electronic Army[12].

Patient Health	Patient EHR
untargeted; may not discriminate	unlikely to target
targeted; may choose specific victims	

[6] http://www.dailydot.com/crime/arrested-atm-heist-45-million/
[7] http://www.eweek.com/security/target-breach-involved-two-stage-cyber-attack-security-reseachers.html
[8] http://www.huffingtonpost.com/2014/09/18/home-depot-hack_n_5845378.html
[9] http://dealbook.nytimes.com/2014/10/02/jpmorgan-discovers-further-cyber-security-issues/
[10] http://www.cnn.com/2015/10/15/politics/malaysian-hacker-isis-military-data/
[11] http://money.cnn.com/2015/10/15/technology/isis-energy-grid/
[12] http://archive.thedailystar.net/beta2/news/new-york-times-twitter-hacked-by-syrian-group/

NATION STATE

These adversaries are the greatest threat likely to be faced. Objectives may be to harm or threaten the harm of one or a group of individuals from an enemy nation, or to obtain the PII and EHRs for targeted or groups of individuals en masse for exploitation. These adversaries have demonstrated extremely high skill and persistence in launching attacks. Notable attacks by these adversaries in other industries are China's Ghostnet campaign to compromise foreign embassy, NGO, news media, and other international organizations[13], North Korean attacks against Sony[14], attacks by Iran against U.S. State Department officials[15], as well as United States and Israeli attacks against Iranian uranium enrichment plants in 2010[16].

Patient Health

untargeted; may not discriminate

targeted; may choose specific victims

Patient EHR

untargeted; may not discriminate

targeted; may choose specific victims

[13] http://www.telegraph.co.uk/news/worldnews/asia/china/5071124/Chinas-global-cyber-espionage-network-GhostNet-penetrates-103-countries.html
[14] http://www.bbc.com/news/world-asia-30670884
[15] http://www.nytimes.com/2015/11/25/world/middleeast/iran-hackers-cyberespionage-state-department-social-media.html
[16] https://www.washingtonpost.com/world/national-security/stuxnet-was-work-of-us-and-israeli-experts-officials-say/2012/06/01/gJQAlnEy6U_story.html

Phase 1

- interview staff
- review network, systems
- inventory, data collection
- review policies, workflow
- **construct attack model**

Phase 2

- design real-world attacks
- ensure patient safety
- launch attacks
- gather results
- **demonstrate empirical threat**

Phase 3

- review attack findings
- discuss with afflicted parties
- design mitigation strategies
- **construct blueprint**

Methodology

Our approach to this research is designed to determine the feasibility of realistic, advanced attacks against patient health in actual hospital settings. Too often research is limited to the specific attack surfaces (e.g., a medical device, a web portal, or a particular software application), and does not demonstrate the full spectrum of the attack possibilities. Because of this, attacks are sometimes deemed unrealistic or too difficult to be practical. **Our research demonstrates that remote adversaries can easily deploy attacks that manipulate records or devices in order to fully compromise patient health**. All research was performed in a *whitebox* setting, meaning all IT staff were fully aware of the experiments and provided certain details to ensure that results were legitimate and that no damage was caused.

The preliminary phase of our research was to collect a wide range of data for which our attacks would be derived. First, we interviewed hospital staff, from the IT department, physicians and nurses, to Biomed departments and some vendors. Next, we reviewed hospital network architectures, network device configurations, critical system configurations, and other high-level design items. Following this, we reviewed the hospitals' processes and procedures that have the potential to affect patient records or health. Next, we assessed hospital policies for security relevant topics, e.g., bring your own device (BYOD), wireless access, remote physician and remote vendor access, etc.

The second phase of our research was to design empirical attacks based on the actual hospital networks, systems, policies, and procedures that we investigated. Attacks were designed to not interfere with actual patient health or records, but to simulate such attacks. For example, we would test attacks on a medical device while disconnected from the network, and in a subsequent step verify that we could access the same versions of said medical device from the network (but not actually perform the attack). Details are given for each attack scenario later in this report.

Attacks were intended to replicate real world attack scenarios as best as possible without interfering with actual patients or records in a way that would actively disrupt day-to-day operations or cause harm. When applicable, attacks were walked through with system administrators, physicians, surgeons, and compliance experts, to determine the real-world ramifications of such attacks.

As part of mitigation, whenever a vulnerability was found, we disclosed all information to the supervisory parties, i.e., the hospital IT or Biomed departments, medical device manufacturers, software providers, or vendors. We worked with those parties to create mitigations for the vulnerabilities found, although a complete, all-encompassing security review of every component was not performed. We advised the affected parties on methods and plans for long-term mitigation strategies, and designed our blueprint strategy around these discussions.

Our focus was on determining attack feasibility and damage from the point of view of compromising patient health or patient records. We did not focus on the specific compliance with regulatory statutes such as HIPAA or HITECH.

Related Work

Attacks that target patient health have been suggested to be possible before, though this research focuses on the exploitation of end-system medical devices that could cause harm. To our knowledge, no real-world attacks have been reported targeting patient health. Research has shown that medical devices are susceptible to compromise, such as pacemakers[17], and insulin pumps[18,19]. Similar attacks have even been demonstrated on simulated patients in a laboratory setting[20]. Though attacks against these systems have only been performed in a research setting, they demonstrate a grave problem. When these or similar attacks are finally exploited in the wild, lives will be lost. In 2015, attacks were documented using medical devices as the pivot onto the hospital's production network[21]. The device was not targeted, but was used to make the attack.

There have, however, time and again been failures of medical devices that have compromised patient health. This report by The Citizen[22] describes numerous failures that resulted in injury. Another report[23] describes 24% of all surgical errors as being equipment related, such as loss of device availability, improper device configuration, and device malfunction. These failures support our hypothesis that attacks that target patient health are viable. If failures can cause harm, and attacks can cause failures, it follows that attacks can cause harm. In fact, it is reasonable to see that targeted, malicious attacks designed to cause failures can do so in non-random, deceptive ways, making them even more difficult to detect and respond to before damage is caused.

Attacks to obtain patient records are prevalent in the media, and on the rise. Highly publicized attacks against Anthem, Tricare, and Community Health Systems[24] show that the spotlight is certainly on this industry at present. Statistics also support this, showing an increase in attacks designed to compromise patient records by 600% in 2014 alone[25]. These types of attacks do not necessarily align with our discussion of attacks against patient health, though one can easily surmise that rampant attacks against a healthcare infrastructure in which patients are actively receiving treatment could likely result in a disruption of that care.

In the past decade, we've seen the emergence of a series of related regulatory statutes through HIPAA, HITECH, and the FDA. These statutes are meant to protect hospital operations, and focus largely on the protection of the privacy and confidentiality of patient health records. These measures have attempted to better protect consumer/patient privacy by creating guidelines, then enforcing them with fines and the aspect of public shame. These statutes have not been successful in curtailing the rise of successful attacks aimed at compromising patient records, as can be seen in the year over year increase in successful attacks. This is no surprise however, since compliance rarely succeeds at addressing anything more than the lowest bar of adversary faced, and so long as more and better adversaries come on to the scene, these attempts will continue to fail.

Lastly, there is wide-spread evidence that advanced persistent threats (APTs) exist and operate within our corporate and government

[17] https://www.umass.edu/newsoffice/article/how-much-security-do-you-expect-your-pacemaker-umass-amherst-expert-works-provide-cyber
[18] http://www.theregister.co.uk/2011/10/27/fatal_insulin_pump_attack/
[19] https://media.blackhat.com/bh-us-11/Radcliffe/BH_US_11_Radcliffe_Hacking_Medical_Devices_WP.pdf
[20] http://www.computerworld.com/article/2981527/cybercrime-hacking/researchers-hack-a-pacemaker-kill-a-man-nequin.html
[21] http://www.computerworld.com/article/2932371/cybercrime-hacking/medjack-hackers-hijacking-medical-devices-to-create-backdoors-in-hospital-networks.html
[22] https://www.citizen.org/documents/substantially-unsafe-medical-device-report.pdf
[23] https://www.citizen.org/documents/substantially-unsafe-medical-device-report.pdf
[24] http://www.modernhealthcare.com/article/20150210/blog/302109995
[25] http://health.economictimes.indiatimes.com/news/health-it/340-increase-in-cyber-attacks-in-healthcare-industry/49111026

infrastructures. Sophisticated attacks have been shown in many industries, including financial[26], media and entertainment[27,28], government[29], education[30,31], social media[32], ecommerce[33], and the list goes on. Even healthcare, which has been demonstrating the prevalence of unsophisticated attacks for years, is now starting to show that advanced attacks are also in the space[34,35]. As it has proven unsuccessful to eradicate these adversaries from other industries, we should approach the problem with the same reasoning that they are in healthcare to stay as well.

Understanding Attacks: Patient Health vs. Patient Records

Fundamentally, the motivations for seeking to compromise patient health vs. patient record assets are very different. On their face, one is meant to cause physical harm and the other is meant to achieve financial gain (with a few exceptions, such as to terrorize or violate privacy). Digging deeper, it becomes apparent that the attack structures and intermediate objectives are very different as well. That is, depending on the attack goals, how the attack is carried out and the resources used will vary greatly. Thus, the defenses against those attacks must also vary. Understanding that there is not a one-size-fits-all solution to infrastructure security is crucial in developing a sound defensive strategy.

Given the below cases, one can quickly see that a staunch focus on protecting PII does not necessarily lend itself to protecting the medical information or the patient, nor does focusing just on the protection of patients or the medical sensitivity. In fact, all three of these motivations should be considered when building a defensive strategy; assuming it really is the goal to protect patient PII, PHI, and patient health.

Targeting PII	Targeting PHI	Targeting Patient Health
Attacks to obtain patient records are most typically untargeted attacks aimed at obtaining personally identifiable information (PII), and not sensitive medical information (personal health information: PHI). The PII is where the value lies. The adversary could care less about the medical situations afflicting the victims of the theft. For the most part, untargeted medical information has no value on the black market.	Targeted attacks to obtain patient records are entirely different. Given the diligence and focus required to target specific individuals' health records, it is likely that these adversaries are capable of obtaining the typical PII found in a medical record by other means. Instead, the goal is actually to obtain the medical information itself. This PHI may exist in many different forms and in many places not necessarily associated with PII, but still linkable to a specific patient.	Attacks against patient health, whether targeted or not, will rarely care about targeting the PII aspects of medical records. Instead, the devices, infrastructure, and specific medical information relating to a patient will be targeted.

[26] http://www.usatoday.com/story/tech/2015/02/15/hackers-steal-billion-in-banking-breach/23464913/
[27] http://www.huffingtonpost.com/2011/04/26/playstation-network-hacker-stole-user-data_n_854106.html
[28] http://www.bbc.com/news/world-asia-30670884
[29] http://www.nytimes.com/2015/07/10/us/office-of-personnel-management-hackers-got-data-of-millions.html?_r=0
[30] http://www.stanforddaily.com/2013/09/23/online-security-breach-prompts-further-security-measures-amidst-uncertain-details/
[31] https://www.washingtonpost.com/local/college-park-shady-grove-campuses-affected-by-university-of-maryland-security-breach/2014/02/19/ce438108-99bd-11e3-80ac-63a8ba7f7942_story.html
[32] http://www.cnn.com/2014/01/01/tech/social-media/snapchat-hack/
[33] http://www.bloomberg.com/bw/articles/2014-03-13/target-missed-alarms-in-epic-hack-of-credit-card-data
[34] http://fortune.com/2015/02/05/anthem-suffers-hack/
[35] http://www.chs.net/media-notice/

Motivations aside, these technologies, medical records, and PII should be protected in lock-step with one another. With few exceptions, nearly all cyber-attacks will leverage the hospitals' infrastructure. To best defend assets within an infrastructure, one must first understand the attacks. To our knowledge, until now there has not existed a comprehensive attack model targeting hospital patient health. After studying hospital workflows, we present the following Patient Health Attack Model that shows how patients are most likely to be targeted in a cyber-attack.

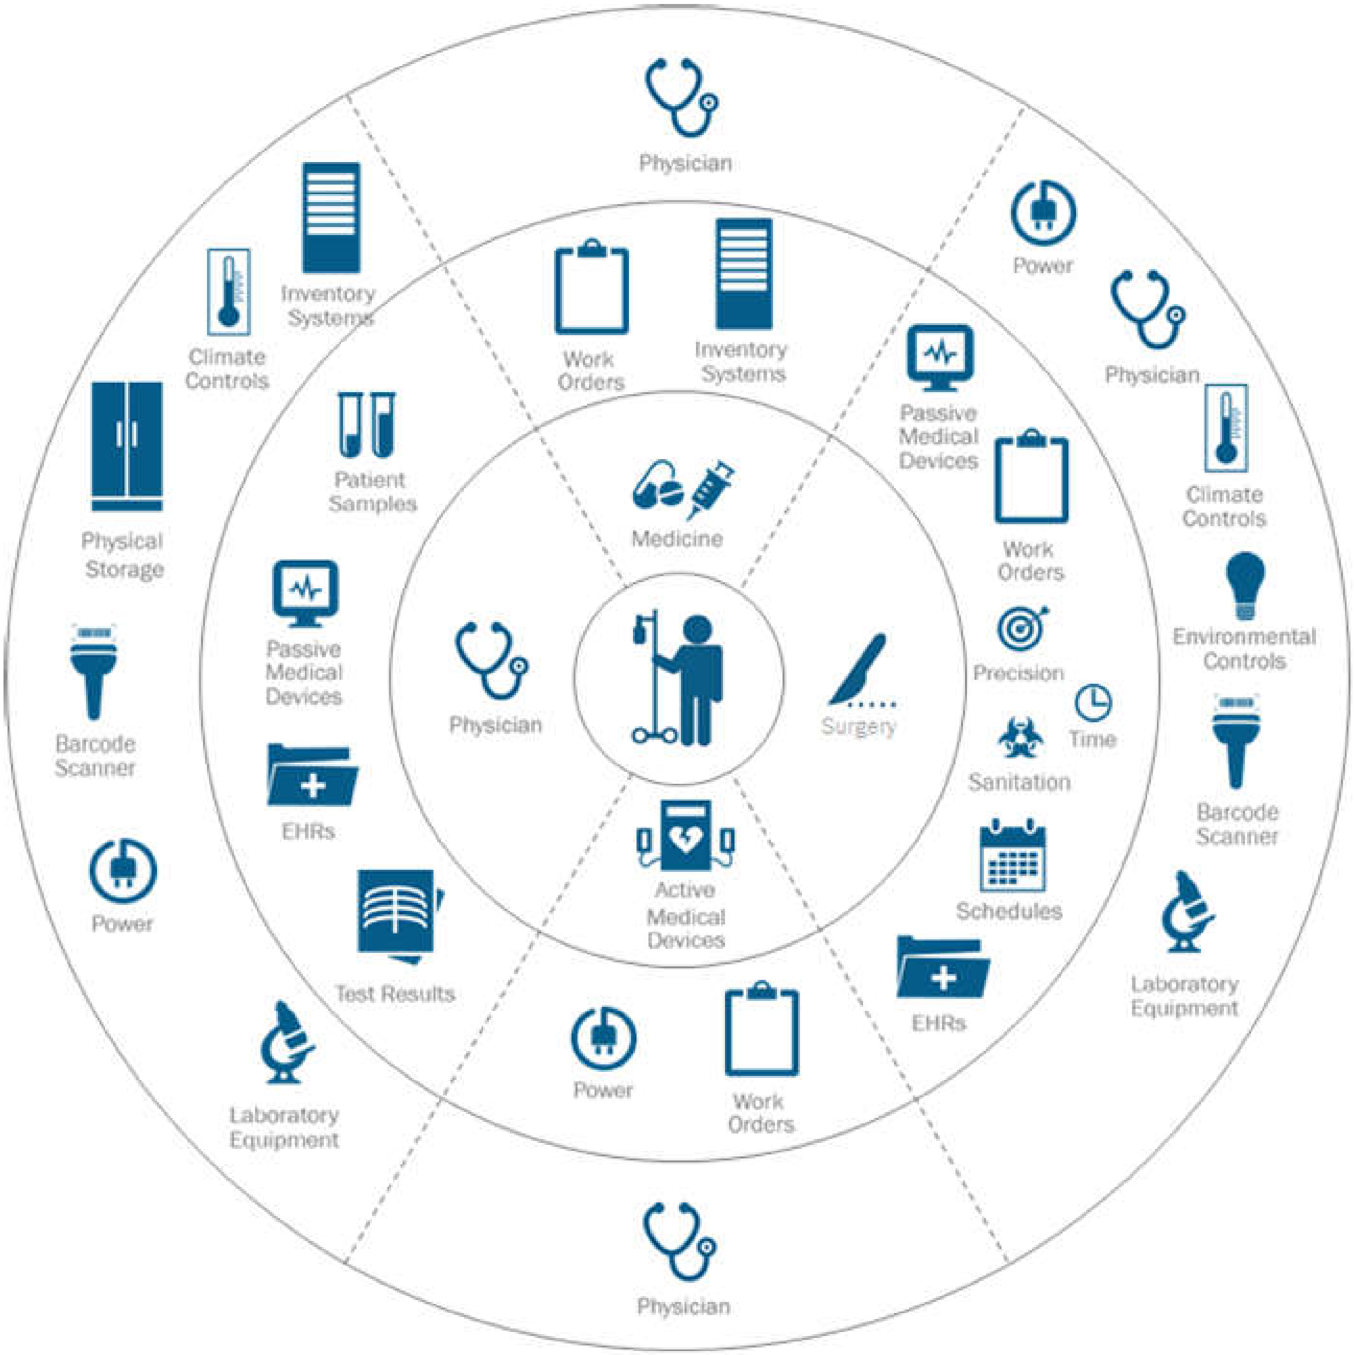

Part II: Research and Results

Patient Health Attack Model

To our knowledge, no comprehensive attack model treating patient health as the target within a healthcare facility has been presented. Our goal in doing so is to help healthcare facilities and security professionals better understand the types of attacks that could be possible that could result in harm to a patient. In the diagram presented, the patient is at the center with attack surfaces that could harm that patient spiraling outward. Primary attack surfaces are those things within a healthcare facility that, if compromised, could directly affect the patient. The diagram then moves outward to secondary and tertiary attack surfaces. There are certainly attack surfaces even further removed from this model, but they have been omitted for brevity. We hope to update this attack model in the future as new attack surfaces are introduced and the overall system evolves, and welcome input that can help us present a more comprehensive list of attack surface classifications should there be any that we've not included.

It is tempting to include in this classification networking equipment, servers, applications, and software, however, these things are not necessarily related to the direct application of the practice of medicine, so they are not included. Certainly, the compromise of a server that contains medical information is an important step in an adversary's attack campaign, however, the server itself is not the attack surface that affects the patient health –it is the EHRs that may be on that system that are part of this classification. In other words, we have intentionally omitted the infrastructure components that are part of an attack campaign, but not part of the administration of care.

PRIMARY ATTACK SURFACES

These are the attacks and attack surfaces that directly affect the patient. That is, if you can compromise one of these devices, it may directly harm the patient as it interacts with them. For instance, controlling an active medical device to deliver a lethal dose of medicine or electricity is a primary attack surface as this touches the patient, whereas altering a medical record is only a secondary attack surface as it requires a physician, or other party to act on the altered information before harm is caused to the patient. Primary attack surfaces are the most crucial to secure.

Active medical devices (AMD) are those devices that interface directly with a patient and administer some medical treatment, which in the event of a compromise could adversely affect the patient's health. These include insulin pumps, heart defibrillators, machines that emit radiation, or any equipment that sustains life, etc. AMDs can be affected to cause harm in the following situations:

- *By denying treatment*. If a device is modified to fail to deliver the necessary treatment, a patient could die or suffer other injury. For example, a heart defibrillator modified or disabled cannot deliver the necessary electrical current to save a patient in distress.

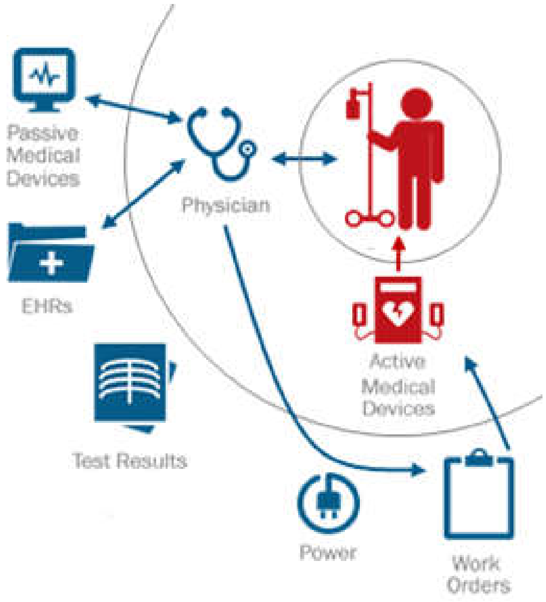

The above diagram shows a compromised AMD modified to cause harm, perhaps by delivering an electrical shock to a patient. Even when all information communicated is good, this diagram shows how a compromised AMD can still directly harm a patient.

- *By modifying treatment.* If a device has been modified to deliver the incorrect medicine, or incorrect dosage, a patient could die or suffer injury. For example, an insulin pump modified to deliver 1000 units or 10mL to a patient could cause death.

- *By modifying to cause harm.* If a device has been modified to cause harm, such as delivering an electrical shock or emitting radiation, a patient could die or suffer injury. For example, an X-ray machine modified to repeatedly emit high levels of radiation could cause harm to nearby patients and physicians.

Medicine and biomaterial is administered to patients regularly in hospital settings, and if the process is compromised this could adversely affect the patient's health. Treatment can be denied, or incorrect treatment administered in the following situations:

- *By altering medical records*. If medical records are altered information necessary to make treatment decisions can be compromised, leading to harmful medical treatment. For example, if a patient with an allergy to penicillin has such a fact removed from their medical records, a physician may prescribe and administer the antibiotic resulting in harm to the patient.

- *By altering work orders. If* work orders instructing nurses or other physicians to administer medicine can be altered, it can cause a patient to die or suffer injury. For example, altering an instruction to deliver X to a patient instead of Y, or morphine to Patient A instead of Patient B, could have catastrophic consequences.

- *By losing or destroying medicine.* The destruction or loss of medicine could affect individual patients or groups of patients by denying needed treatment. This could happen if climate control systems are disabled, spoiling medicine, or if work orders are issued to exhaust a supply.

- *By altering medicine inventory.* If medicine inventory can be altered, unsuspecting physicians, nurses, or pharmacists could treat a patient with the wrong medicine. Clearly, this could have adverse effects on a patient.

- *By altering the transport of medicine.* If medicine cannot reach a patient, the wrong medicine reaches a patient, or the medicine reaches the wrong patient, it could cause harm.

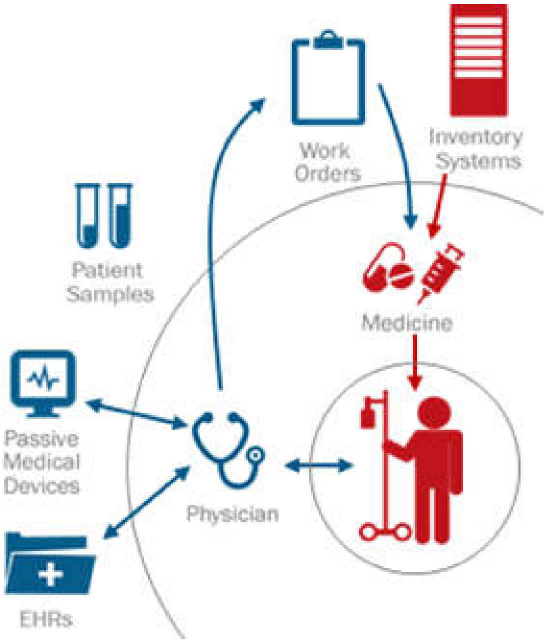

The above diagram shows how an attack that disrupt the accurate delivery of medicine can directly affect a patient's health. In this case, altering a medicine dispensary's inventory to produce the wrong medicine, or wrong dosage.

Surgery and procedures are administered to patients regularly in hospital settings, and if the process is compromised could adversely affect the patient's health. Surgery or treatment can be denied or mis-administered in any of the following situations:

- *By altering work orders.* If the directive for work is altered, even slightly, it could cause a patient to undergo incorrect surgical procedures. For example, a misdirected surgeon could amputate the wrong leg, or remove organs from the wrong patient.

- *By altering medical records.* If medical records are altered, information necessary to perform the surgery correctly could be missing or erroneous, leading to a harmful surgery. For instance, if a patient's blood type is changed, and a transfusion is required during surgery, the wrong blood type could be used causing the patient to have a serious reaction and possible death. Also, altering imaging studies or falsely switching or relabeling imaging studies could lead to gross surgical or interventional errors. Even so much as altering patient vitals, such as height, weight, or age could affect an anesthesiologist's ability to perform his or her job.

- *By altering surgery schedules.* If a surgery schedule can be altered, it could cause an imperative surgery to be postponed, denying service to the patient and possibly causing harm.

- *By denying or interrupting remote access.* It is common today for surgical procedures to be conducted by physicians remotely. Disrupting this service could cause harm to the patient undergoing the procedure. For example, interruption of communications between an onsite surgeon and distant advisor/consultant could result in a failed procedure, compromising the outcome or even the patient's survival. This problem could become more significant as we see the growth of telemedicine and teleconsultation.

- *By compromising surgical precision.* If medical and support equipment can be made to fail during a surgery, this could negatively affect a surgeons ability to perform the surgery without error. For example, if a heart or blood pressure monitor could be made to fail during surgery, or to report inaccurate readings, this could lead a surgeon or anesthesiologist to make a mistake or overlook key patient vitals during surgery, resulting in patient harm.

- *By compromising surgical expedience.* During surgical procedures, time is of the essence. Attacks that inject delay in to a surgical procedure could result in harm to a patient. For example, problems related to primary functions in the surgical theatre, such as lighting, anesthetic equipment, failure of intraoperative imaging (ultrasound, X-ray, CT, etc.) and failure of basic surgical tools such as electrocautery and suction could sufficiently lower the patient's chances of survival.

- *By denying organs, blood, or other vital biological components.* Surgeries often require the infusion of blood, or the transfer or replacement of organs. If organs required for surgery can be maliciously lost, damaged, or destroyed, it could result in a failed surgical procedure. For example, if the climate control system for storing necessary blood or organs is disrupted, damaging the biological material, this could result in a patient failing to receive proper surgery.

- *By disrupting the environment.* Surgery requires a sanitary, well-lit, peaceful, cool, and restricted environment. If the environment where a surgery is being performed can be disrupted, it could cause surgeons to make mistakes. For example, an

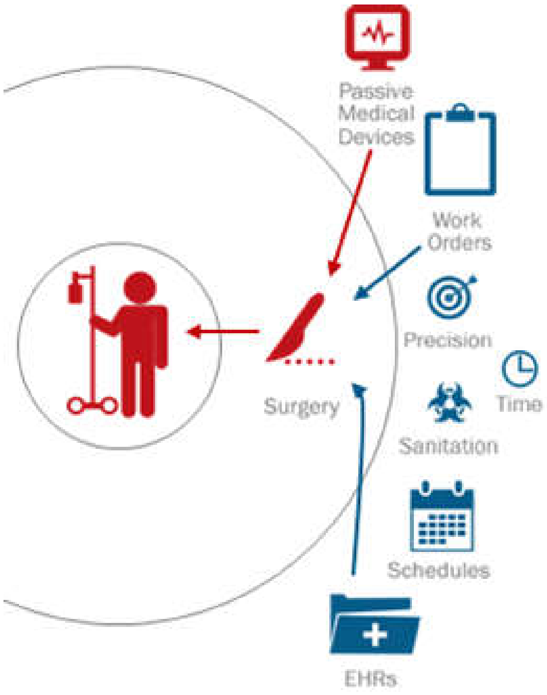

There are many attack surfaces that can disrupt a surgery. In the above example, a compromised monitoring device is made to display incorrect information about the patient, hindering the surgeon's performance.

attack that could disable lighting, or set off a fire alarm or sprinklers, could blind or deafen a surgeon or disable the environmental controls related to heating or cooling, likely resulting in non-optimal surgical precision.

Clinicians are those hospital staff who directly care for and treat patients, and if these clinicians can be manipulated it could adversely affect patient health. Clinicians can be misdirected, misinformed, and even distracted to the point that caring for patients is denied, as discussed in the following situations:

- *By misinforming clinicians.* If the information provided to a clinician is incorrect, such as a patient monitor indicating a heart related emergency, hospital staff could react appropriate to the technology indicators, but not appropriately to the actual situation. If the result is that the patient is shocked with a defibrillator or treated unnecessarily with anti-arrhythmic drugs it could harm that patient.

- *By misdirecting clinicians.* If hospital technology can be made to misdirect or distract clinical staff, this could result in patients not receiving needed emergency care. For instance, issuing alarms in parts of the hospital where there are no events occurring could draw staff away from patients who are in actual need of attention.

SECONDARY ATTACK SURFACES

These are attacks or attack surfaces that don't harm patients directly, but support attacks against primary attack surfaces. That is, each of these attack surfaces, if compromised, still requires an additional step before a patient is harmed. Even so, this should not be taken lightly. Secondary attack surfaces are very likely to affect a patient if compromised.

Passive medical devices (PMD) are devices that report on patient vitals or other information needed to inform or alert clinical staff of medical events or needed treatment. If these devices are compromised, they could affect patient health through a clinician in the following ways:

- *By reporting false information.* If incorrect information is reported to clinical staff, they are more likely to make life threatening decisions with regard to patient treatment, such as not addressing a condition or applying the wrong treatment.

- *By not reporting medical events.* If PMDs do not report medical events, such as a heart attack, this could prevent that patient from receiving timely medical treatment for the medical event.

- *By reporting false medical events.* If a false medical event is reported, a patient could receive relatively damaging treatment to the true state of their health, such as being shocked with a defibrillator when there is no legitimate heart event occurring.

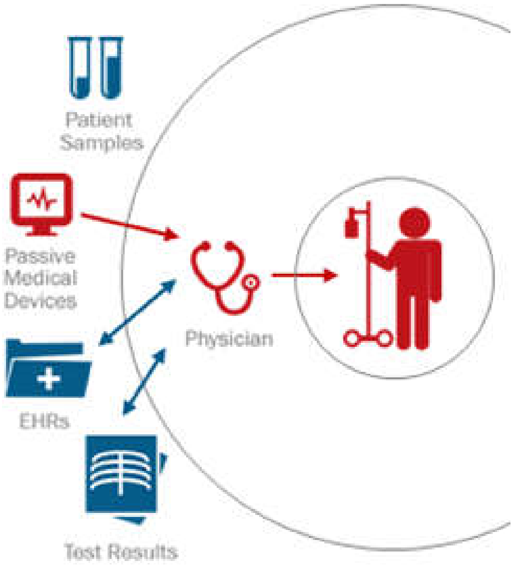

In the above example, a patient vitals monitor sounds and alarm indicating the patient is having a heart attack, causing physicians and nurses to rush to the patient's aid and apply unneeded medical treatment.

Electronic health records (EHR), often the target of attack when the goal is identity theft or fraud, could be modified in a way that leads to patients receiving inappropriate or harmful medical care from a clinician or surgeon:

- *By containing false information.* If incorrect information is reported to clinical staff, they are more likely to make life threatening decisions with regard to patient treatment, such as not addressing a condition or applying the wrong treatment.

- *By failing to supply information.* If incorrect information is reported to clinical staff when immediately relevant, like during a surgical procedure, this denial of service could have a direct negative affect on the outcome of the procedure.

Medicine inventory systems are databases or physical storage that tell exactly which medicine is which. If compromised, medicine dispensary systems and clinical staff could be provided with incorrect medicine or dosages that affects patients in the following ways:

- *By delivering incorrect medicine.* If a medicine inventory system swaps the names or designations of medicine, a mistake could occur whereby a clinician obtains and treats a patient with the wrong medicine.

- *By delivering incorrect dosages.* If a medicine inventory system changes the quantity of dosage information for stored medicine, a mistake could occur whereby a clinician obtains and treats a patient with the wrong medicine dosage.

Power is sometimes provided or regulated by connected devices such as uninterruptable power supplies (UPS), generators, or even building infrastructure power in general. It can be disrupted to cause harm to patients in the following ways.

- *By disabling AMDs.* If an active medical device required to monitor and immediately respond to a patient health event, disabling the power to these devices could cause them to fail to deliver this treatment.

- *By disrupting surgery.* In surgery there are a variety of systems that if they should fail could affect the successful outcome of the surgery, such as passive medical device monitors, environmental controls, lighting, and access to systems that provide medical information.

Patient samples are physical matter taken from a patient for the purposes of reaching a diagnosis, such as blood, urine, bone marrow, etc. These materials are generally collected from a patient, transferred, and stored prior to and after testing. If patient samples are compromised, it could adversely (but not directly) affect the patient's health.

- *By corrupting samples.* If patient samples are corrupted, it could lead to an inability to diagnose, or a misdiagnosis. This could happen by altering lab equipment settings or environmental controls meant to maintain sample longevity.

- *By losing or destroying samples.* If patient samples are lost or destroyed, it could lead to an inability to diagnose. This can happen

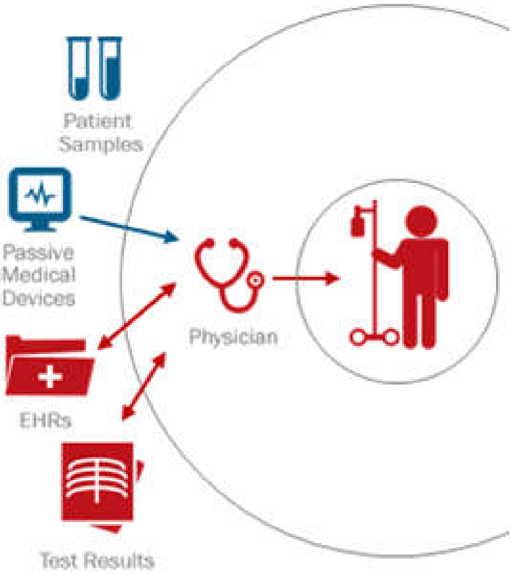

In the above example, EHR integrity is compromised and test results are altered, causing physicians to make incorrect medical decisions that adversely affect the patient's health.

through the alteration of work orders and the transport of samples, or the modification of sample inventories and climate controlled environments.

Test results are produced from any number of laboratory tests or direct patient interaction, such as information reported following a metabolic panel, or the images created by an x-ray. Test results are generally collected from a patient or from patient samples, and recorded for later use. If a patient's test results are compromised, it could adversely (but not directly) affect the patient's health.

- *By corrupting test results.* If patient test results are corrupted, it could lead to an inability to diagnose, or a misdiagnosis. This could happen by altering lab equipment, lab procedures, or stored test results.

- *By modifying test results.* If a patient's test results are maliciously modified with the intent of covertly causing harm, it could lead to a misdiagnosis or other treatment that could harm a patient. This could happen by altering a patient's medical records in such a way that test results are modified, or by compromising test equipment and instructing it to report incorrect results. For example, if a patient's medical records were altered to change the patient's blood type, this could lead to serious complications during surgery, dialysis, or even a general transfusion.

Organs and other high demand material (such as blood or bone marrow) are essential to sustaining the life of some patients. Attacks that could deny access to these materials could prevent a patient from receiving the life-saving treatment they need, or even prevent the successful completion of a surgery depending on the timing of the attack. In doing so, these attacks do not directly affect the patient, but indirectly can have a significant impact on treatment. Denial of organs can occur through a variety of attack scenarios.

- *By altering the transport of organs* from donor, to facility, to patient, etc. If the necessary organs or other materials cannot reach the patient, the patient could die during or while waiting for surgery.

- *By altering climate controlled transport or storage of organs* an attacker could corrupt the organs, thereby denying the patient the treatment they need.

- *By altering transplant recipient lists* an attacker could prevent a patient from receiving the medical response they need as other parties instead receive the organs necessary for survival.

- *By altering the medical information of an organ donor* an attacker could invalidate the organs for use by anyone, or just certain parties in a targeted scenario.

- *By altering the medical records of the candidate* recipient an attacker could invalidate that patient's candidacy, preventing them from getting the treatment needed for survival.

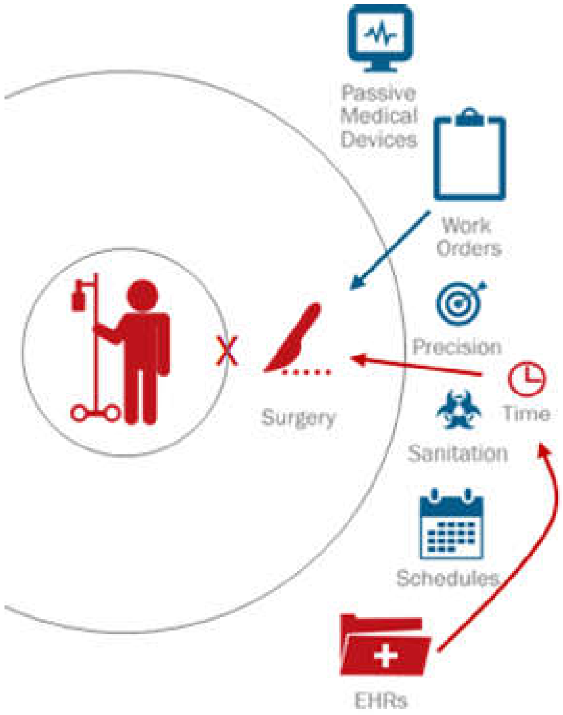

The above diagram shows an attack on an organ transplant list; a secondary attack surface. By removing or invalidating a candidate in need of organs, it does not directly harm the patient, but if the problem is not addressed that patient could still suffer greatly from the attack.

Fear, uncertainty, doubt. By use of any or all of these malicious techniques a coordinated attack by a single individual, group, entity, or government could create widespread fear and concern regarding the safety of the medical and health system in general. The result could be far reaching in that it might cause patients needing medications, medical or surgical attention to delay or refuse treatment due to reports of widespread harm to individuals caused by hackers.

TERTIARY ATTACK SURFACES AND BEYOND

We don't delve into details of attack surfaces further than two degrees of separation from the patient. There are far too many possibilities. Everything from financial systems, clinical staff workstations, patient portals, kiosks, and mobile devices are all attack surfaces that can be leveraged to launch attacks against a hospital. The use of these types of attack surfaces is detailed in our attack anatomies section of this report. In general, a hospital should address the security concerns of all aspects of its infrastructure, but this attack model serves to identify the highest risk areas and the most critical to secure when the target asset is patient health.

Attack Anatomies

This section describes actual attacks demonstrated and emulated by our researchers in actual field settings that would have resulted in patient injury or death if launched by malicious parties with such intent. The purpose of this section is to illustrate actual attack anatomies that we found to be successful, rather than enumerate all full attack scenarios that were possible. In fact, in most cases the same result (injury or death) could be achieved in a multitude of ways beyond what is described in this report.

Our first priority was safety, thus all attacks were performed with the supervision and permission of hospital personnel authorized to perform such experiments. All attacks were performed on either non-critical systems, decommissioned or non-connected medical devices, or other systems in manners that would not affect patient health in any way. In most cases, all but the final step that involved manipulation of an actual medical device, medicine dispensary, or health record was performed online, with the final step taken offline to ensure there was no accidental injury or harm caused to a patient.

External attack to manipulate active medical device

In this attack scenario, we demonstrate that a foreign group could launch an attack against patients in a US hospital, leveraging passive medical devices to cause those patients harm.

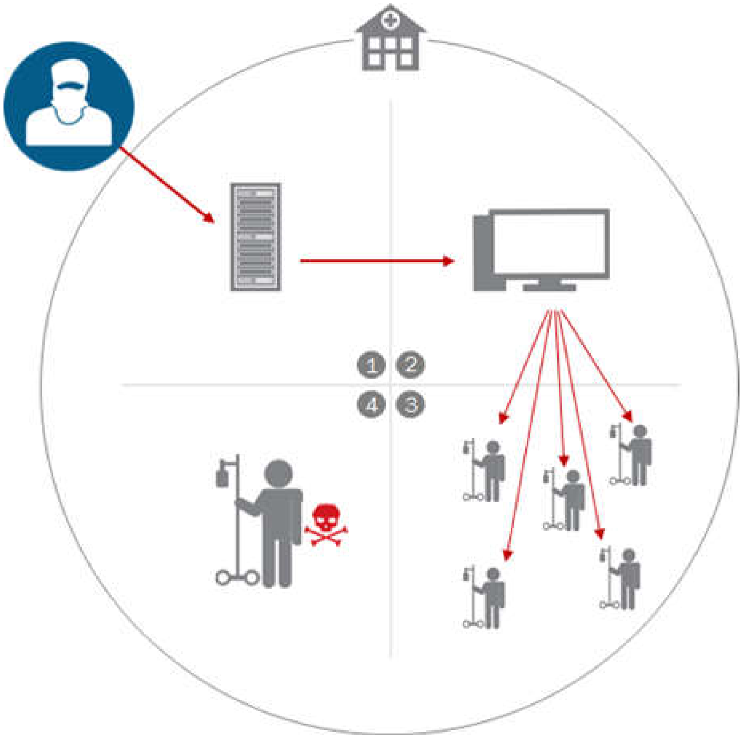

Step 1 – Circumvent the perimeter
Our team targeted an externally facing web server at one of the hospitals. We were able to exploit vulnerabilities in the server to gain control of the web server, thus gaining a foothold on the internal network.

Step 2 – Pivot within network
Once inside, our team was able to perform scans of the internal network without detection and identify systems on another network segment that appeared vulnerable. We identified additional vulnerable systems within the network, thus gaining additional footholds now in more desirable network locations.

Step 3 – Probe the network
From our new vantage points, our team was able to identify numerous passive medical devices (patient monitors) on the network, of which **all** were vulnerable.

Step 4 (offline) – Compromise a medical device
On a disconnected network segment, our team demonstrated an authentication bypass attack to gain access to the patient monitor in question, and instructed it to perform a variety of disruptive tasks, such as sounding false alarms, displaying incorrect patient vitals, and disabling the alarm. This attack would have been possible against all medical devices from step 3, likely preventing assistance and resulting in the death or serious injury patients.

Analysis The above attack scenario is harrowing. Diligently executed, many human lives could be at stake, and extrapolating this problem to other hospitals is even more worrisome. While there are logistical considerations for how a targeted attack on an individual could be launched, it is very clear that random attacks are possible and viable.

Mitigation The above attack would have to be partly mitigated by the following: Application security assessments of all outward facing services, proper network segmentation to support a secure, scalable infrastructure, appropriately deployed intrusion detection systems to detect early signs that an attack is in progress, and air-gapping (or heavily restricting access) to networks with active medical devices as recommended by the blueprint in this report.

Lobby attack to manipulate medicine/bloodwork workflow

In this attack scenario, we demonstrate that from the lobby of the hospital an attacker can manipulate the flow of medicine or blood samples within the hospital, resulting in the delivery of improper medicine types and dosages, as well as the mixing up of blood samples.

Background – some of the hospitals we reviewed use patient, medicine, and sample tracking software, and hardware barcode scanning and label printing devices. These systems are used to more quickly and easily track, verify, and deliver medicine and samples – ideally reducing oversight by requiring a digital check of patient-to-medicine before administering a dose.

Step 1 – Circumvent the perimeter

Our team entered the hospital, and accessed a vendor kiosk in the hospital lobby (prior to passing through security). Our team was able to break out of the device's "kiosk mode" and access the underlying system with full control, thus gaining a foothold on the internal network (the system was not on a restricted network zone).

Step 2 – Pivot within the network

Pivoting within the network was unnecessary, as the kiosk was located on a wireless LAN with access to the desired target end systems.

Step 3 – Compromise the end system

Our team was able to identify numerous mobile computer stations (i.e., the mobile stations found in most emergency and hospital rooms), of which one was readily exploitable.

Step 4 (offline) – Manipulate connected equipment

From the compromised mobile station, we had access to the medicine and bloodwork barcode scanning device. With that access, the adversary is able to view the patient name and identification information (i.e., the barcode values) of the patient in the room, as well as control the barcode scanning device to report correct scans of patient labels when there was in fact a mismatch.

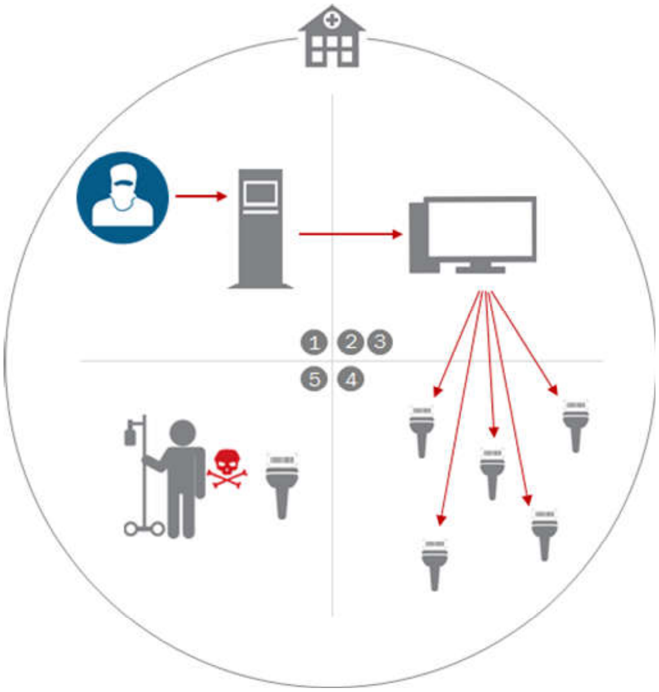

Hypothetically, by compromising two mobile stations in this way, our team could have printed the labels of one patient with the identification number of another, contaminating samples, and possibly causing inappropriate treatment to be administered.

Analysis The above attack scenario emphasizes the phenomena that as technology becomes ubiquitous, attack surfaces increase and security vulnerabilities become prevalent. The modern use of an information kiosk, and sample tracking technology, while both having use within the hospital to better patient experience and reduce clinical errors, they added two distinct attack surfaces that would ordinarily not be present. That being said, the solution is not to bar these systems by any means, but instead to work toward better security for them.

Mitigation The above attack scenario would have been mitigated in a variety of ways, including: Application security assessment of the kiosk device, segmentation of the kiosk device on a non-internal or public network, security hardening of hospital end systems, and application security assessment of the sample tracking servers and appliances.

EHR system compromise to issue improper treatment

In this attack scenario, we demonstrate that without ever targeting a hospital system directly an adversary can cause actions within a hospital to occur that could harm or kill a patient. This attack targets an EHR platform and its users, neither of which are on site at a hospital, but for which information entered is used at the hospital.

Step 1 – Construct XSS attack

The objective of the attack being to instruct an EHR web application to perform an action on the behalf of a physician. One EHR system reviewed as part of our research that was vulnerable to a variety of cross-site scripting (XSS) attacks. These types of attacks are common (found on the OWASP top 10 list of web application vulnerabilities) and readily exploitable. The XSS attacks identified allowed for the modification of administrator settings, the addition of users, and thereafter the direct manipulation of patient records.

Step 2 – Construct an attack payload

If a payload could be delivered and executed by an administrator, it would elevate our privileges. We constructed a payload that could be delivered through an unprivileged nurse or physician account that would escalate our privileges to that of an application administrator. This was done by entering basic patient information about a fake patient with a low privilege account, creating a persistent cross-site scripting payload. When the administrator would review activity for the system, the payload would be executed by the administrator elevating our user account privileges.

Step 3 – Launch the attack

In a test environment, we performed the attack by entering the payload in to the patient information fields. We then demonstrated that these attack payloads would automatically execute when viewed by an administrator.

Step 4 – Log in and manipulate records

Once our user account privileges were escalated to that of an administrator, we logged in with the full ability to modify the health records of all patients in the database. Modification of the health records is a viable means for creating a situation that could harm or kill a patient.

Analysis The above attack scenario illustrates the extensive field of the healthcare ecosystem and that attack surfaces beyond a hospital can still result in harm caused within the hospital walls. This notion is common in most industries as cloud services have risen to meet the technological needs of the marketplace. It is crucial to realize that securing the healthcare ecosystem involves many parties, technologies, and services.

Mitigation The above attack could have been mitigated by the following: Application security assessment of the EHR system, training of the physicians and possibly training of the EHR system developers, and a variety of application-level defense in depth measures if built into the EHR platform.

USB stick used to gain network foothold and manipulate medicine distribution

In this attack, we demonstrate that without ever touching a computer system on the hospital network (and arguably without ever setting foot in a hospital), an attacker can gain a network foothold from which he can then attack critical medicine dispensary equipment, causing the improper medicine to be dispensed, resulting in patient harm or death.

Step 1 – Prepare malware infected USB sticks

Our team prepared 18 USB sticks, each infected with simulated malware (created by us and tested to be benign) that emulate the download and installation of malicious software designed to take control of the target end system and allow a remote adversary to then remotely control that system. When the malware infected USB sticks are plugged into a systems USB port, the machine is typically infected.

Step 2 – Deliver malware infected USB sticks

We entered one of the hospitals and floor-by-floor planted the 18 USB sticks in areas likely to be discovered and installed by hospital staff. The USB sticks were marked with the logo of the hospital.

Step 3 – Wait for infection

Within 24 hours of planting the USB sticks, the sticks were used at nursing stations which requested malware from our server.

The nurse and physician operated devices represent a special point of failure, as not only does the adversary then control a machine on the internal hospital network, but can harvest the credentials and access the accounts of the nurses and physicians who use those machines.

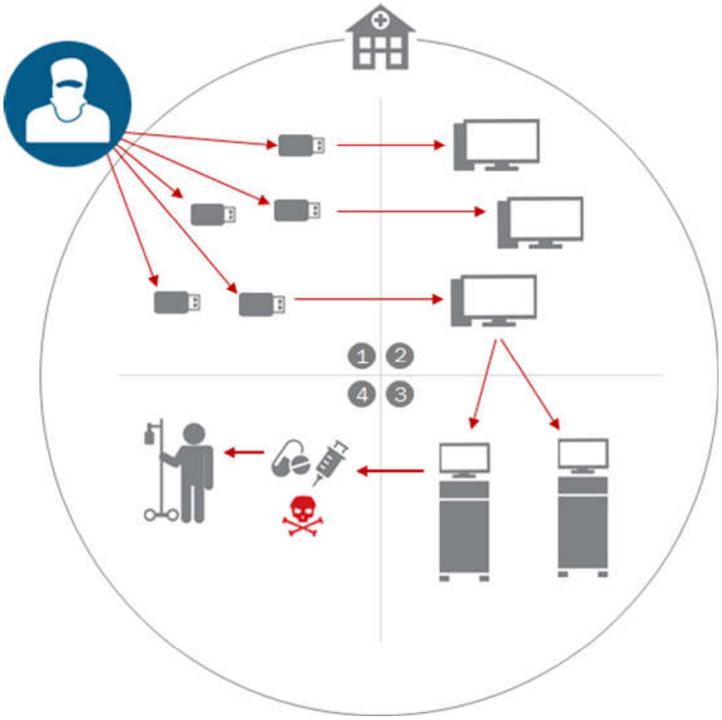

Step 4 – Identify a medicine dispensary device

From the compromised machines, we confirmed that we had numerous vantage points from which to attack numerous medicine inventory and dispensing systems.

Step 5 (offline) – Compromise medical dispensary device

At the time of this reporting, we are working to demonstrate that an attack against the particular dispensary is possible, meaning that anyone who can connect to the dispensary can then get access to the configuration interface and manipulate what the device believes it has to be its inventory. If this medication were then given to a patient, it would likely harm or kill the patient.

Analysis The above attack shows that sophisticated adversaries can leverage low-tech means to bypass certain security defenses, in this case, perimeter defenses that protect internal systems. The scenario also demonstrates that there are multiple avenues for performing any particular attack.

Mitigation The above attack could have been mitigated by the following: endpoint security controls preventing unknown devices from being connected, training of staff to understand the threat, segregation of the network as to keep critical systems such as the medicine dispensary on a separate network zone as the likely to be infected end systems, and application security assessment of the medicine dispensary devices.

Many more scenarios

The examples listed above represent a small fraction of the attack scenario possibilities that could result in the injury or death of a hospital patient. Furthermore, these attacks were based on security assessments of relatively few hospitals compared to the greater healthcare ecosystem which incorporates thousands of such facilities. We draw from these facts that cyber-attacks aimed at harming a patient are not only feasible, but that typical hospital security postures are not appropriate for defense against this type of adversary. In the following sections, we discuss the most prominent security shortcomings found across most hospitals investigated.

General design issues with hospital security

Security design issues are those decisions made by the organization which, if followed, should best protect the organization's interests. For hospitals these interests include (or should include) protecting patient health, and then protecting patient records and other hospital assets. These high-level, organizational or department decisions affect the manner in which boots-on-the-ground employees of the organization may implement specific security measures.

The following are the variety of security design issues we found with the hospitals investigated. These issues in part explain the vulnerabilities we discovered, but more importantly they point to a larger problem: that hospitals cannot simply be *fixed* by patching systems and bolting on defenses. The manner in which security is understood at these sites must fundamentally change before effective security can be implemented.

Lack of funding

Enterprise focus: Business

Arguably the most detrimental issue with hospital security is the lack of funding available to both design and implement good security. Most other issues stem from the fact that the funding is simply not available for these purposes. The issues aren't so much that hospitals do not have the funds, but that they are directed in a way that security is not a priority. This needs to change in order to protect patient health.

The average hospital has a comparatively low budget allocation for information security in relation to other industries, which was confirmed by our research. We found that waste was also an important factor in hospital budgets. Much of the available funds were spent on lower priority or less effective security measures, or appropriate security measures that could not be effective due to other shortcomings, such as a lack of staffing or other equipment.

Until hospitals can devote the appropriate funding toward their infrastructure's security, it is unlikely that much security progress will be made. Certain minimums must be met before any reasonable efforts can exist. For instance, hiring a minimal staff to be able to implement, or even just oversee/organize the implementation of security measures by others is necessary before anything of value can occur.

Our blueprint given in this report describes the appropriate hospital spending per category in order to achieve a reasonably strong security posture.

Lack of appropriate staffing

Enterprise focus: Business

Without people, implementing security is not possible. We found that most hospitals had between 0 and 1 dedicated information security staff. That is to say that a variety of staff do participate in the security process –most hospitals had numerous information technology staff, physical security guards, and even upper management personnel with a directive to ensure the security of the infrastructure– but there were typically between 0 and 1 staff who were both wholly and specifically directed to maintain the infrastructure security and had the requisite training or experience to take this on.

We found that most security functions were performed by information technology staff who did not have the advanced training necessary to be effective, and likewise the staff with the knowledge were in director positions and did not have the individual resources to address security issues as well as all other departmental needs.

Any hospital with greater than 20 beds should employ at least one fulltime security professional directed to provide for the security of the infrastructure as their chief directive. Larger facilities should maintain an IS staff-to-bed ratio of at least 1:75.

Our blueprint given in this report describes the appropriate staffing allocation for security based on hospital type and size necessary to achieve a reasonably strong security posture.

Information Technology Staffing*	
Healthcare	Others
~5	~10

* Staff per one hundred other staff.

Information Security Staffing*	
Healthcare	Others
0-1	~2

* Staff per one hundred other staff.

Lack of effective training

Enterprise focus: Business

We found that most hospitals implemented no (or minimal) security training. It is important that all staff throughout the organization receive some level of information security training. Clearly, information technology and help-desk staff must be trained to understand security fundamentals and specifics alike, as they are the most likely to come across system misconfigurations, attack evidence, or malware infections. Additionally, nurses, physicians, technicians, and other non-IT staff should undergo security training to broaden their awareness, help them to avoid common mistakes, and to help them to detect attacks and misbehavior on their own, all contributing to the greater security posture of the hospital. Even the decision making, C-level, board member, and other management staff should undergo relevant training. Unless they understand the reality of the threat landscape, the ramifications of changing technology, and the proper way to design and budget for security initiatives, the organization will have great difficulty staying ahead of what is needed to protect patients and their privacy.

All hospitals should administer or outsource the administration of role-based security training for all hospital staff who may interact with technology, or be responsible for those departments that control security or technology. Depending on the employee’s role, the organization should anticipate all such staff receiving training on a recurring basis of between once per year for typical staff and quarterly for specialized staff, anywhere between 4 and 40 hours per year.

Our blueprint given in this report describes the appropriate training requirements and hospital budget required for meeting these needs.

Improper organizational structure

Enterprise focus: Business

No hospital we investigated separated information security (IS) from information technology (IT). In fact, all IS responsibility fell within IT. Since the IT and IS departments have conflicting directives (openness and functionality for IT, and closedness and restriction for IS) it is inappropriate for one to fall under the purview of the other. In many cases, we experienced the IS and IT decisions being made by the lowest level staff responsible for deploying and configuring the technology, when in fact these decisions are of critical importance and should be made at the department level with direction to separate staff as to how it should be deployed and configured.

The classic example illustrated this needed separation is that of a firewall. An open firewall configuration provides for the functionality of network services without any hindrance, but leaves the network open to attack, while a closed firewall configuration limits network functionality and mitigates potential attacks. While there is probably a balanced configuration that meets both the goals of functionality and security, if the party responsible for making those decisions receives strong pressure for one goal and not the other, they will always gravitate toward meeting that goal at the expense of the other. Since functionality drives revenue and is always the immediate need, it routinely trumps security.

Much like other departments within the organization, IS and IT should check and balance against each other. An organization wouldn't have Accounting fall under Marketing, and vice-versa, because the interests of the overseeing department would take precedence and the lower department's goals would be less likely met. The same is true of IT and IS.

Hospitals should adopt the modern organizational hierarchy placing Information Security (IS) responsibilities in a department separate from Information Technology (IT), but that report to the same executive department upward in the chain, such as the CIO, CEO, or even the board.

The cost of such a change can vary greatly between organizations, but it is an essential step in our blueprint toward reaching a stronger security posture.

Lack of defined, implemented, and/or auditable policy

Enterprise focus: Policies and Procedures

Policy is in part the definition of how security controls should be implemented, but more importantly defines the organization's goals, and the details on how to uphold those goals on a regular basis. Well-articulated security goals and the detailed requirements on how to achieve them should be laid out for the organization just as any other corporate policy. When policies are poorly defined, or not defined at all, they are prone to misinterpretation and the organization is prone to security misconfiguration and failure.

We found that there was a general lack of security policy material at the hospitals we investigated, that policies were not implemented or enforced, or that the requirements of those policies were not verifiable through audit. For instance, most had a written *password* policy, but it was not adhered to in practice. In some cases, there was no defined network policy to compare the implementation against.

Since security starts with the well-defined policy by which security measures should be implemented, it is critical that healthcare organizations direct their focus on these soft topics so that actual implementation may succeed. Without it, it will be unlikely that security can be measured in any significant way.

Lack of network awareness

Enterprise focus: Technical

As organizations grow, holistic awareness of the organization becomes more difficult. This is certainly true of the technological aspects of an organization, such as its deployed networks, systems, software, and user base. Without an awareness of these components, security holes develop and it becomes increasingly more difficult to remediate problems as they arise.

We found that few hospitals had a robust awareness of their network infrastructure, or the systems deployed within it. As a result, we found systems that were beyond their end-of-life, unpatched, or misconfigured, some of which could not be easily located physically in order to decommission, though they were still active on the network. It was clear that a lack of awareness was leading directly to security vulnerabilities.

An organization such as a hospital should maintain detailed, up-to-date network, system, software, and user inventories. This aids in the fast detection and remediation of problems, ensures system maintenance is current, and allows policy to be verified against implementation.

Lack of audit procedures

Enterprise focus: Policies and Procedures

With the exception of our investigation, nearly all hospitals we approached performed no recurring audit of their infrastructure to ensure adherence to policy, or that network and system vulnerabilities did not exist. This step is crucial in an organization maintaining a strong security posture. It is important to periodically review not only the deployment of systems to ensure they adhere to policy, but to reassess the policy itself, to ensure it is current with the state of the art. Hospitals can self-audit or engage outside parties to assist in these procedures, and in all likelihood, both are recommended.

Lack of logging/monitoring

Enterprise focus: Technical, Policies and Procedures

Detection of a compromise or an attack in progress is essential to any organization staying astride with the adversary. In practice, not all attacks will be prevented, and so detection of attacks in progress or that have been successfully carried out is paramount to reducing damages. This is most commonly done in two steps: the logging of network and system events and behavior, and the monitoring of that data to assess whether or not abnormal behavior is exhibited and if it should be investigated further.

Of the hospitals we investigated, few incorporated a robust logging strategy, and fewer still monitored those reports to determine if malicious behavior had been captured. These results are also essential for forensic purposes after the fact, but become rather useless when not used correctly.

Insecure network architecture

Enterprise focus: Technical

An organization’s network architecture best suits security when it is designed to separate systems and network zones based on functional business purpose, or the assets that reside within that zone. There are other valid approaches as well, as the circumstances may dictate, but no matter what the long term scalability of security within that network must be considered.

Most of the organizations’ networks that we investigated were not architected with security in mind. As a result, they did not support a strong implementation of security controls, and in some cases significant effort was required in order to restructure those networks to that a more appropriate architecture.

Hospitals in particular must focus on the importance of network architecture, as the assets they have to protect are of great value and involve the incorporation of technologies atypical to most industries, such as active and passive medical devices. Our blueprint found in this report recommends full, heavily restricted segregation of any network where active medical devices are connected to living patients. Similarly, heavy restrictions are placed on networks that protect secondary attack surfaces, such as medicine dispensaries, environmental controls, etc. These architectures are not typically deployed today in the healthcare ecosystem, but it is essential that some form of it be adopted to protect our most precious assets.

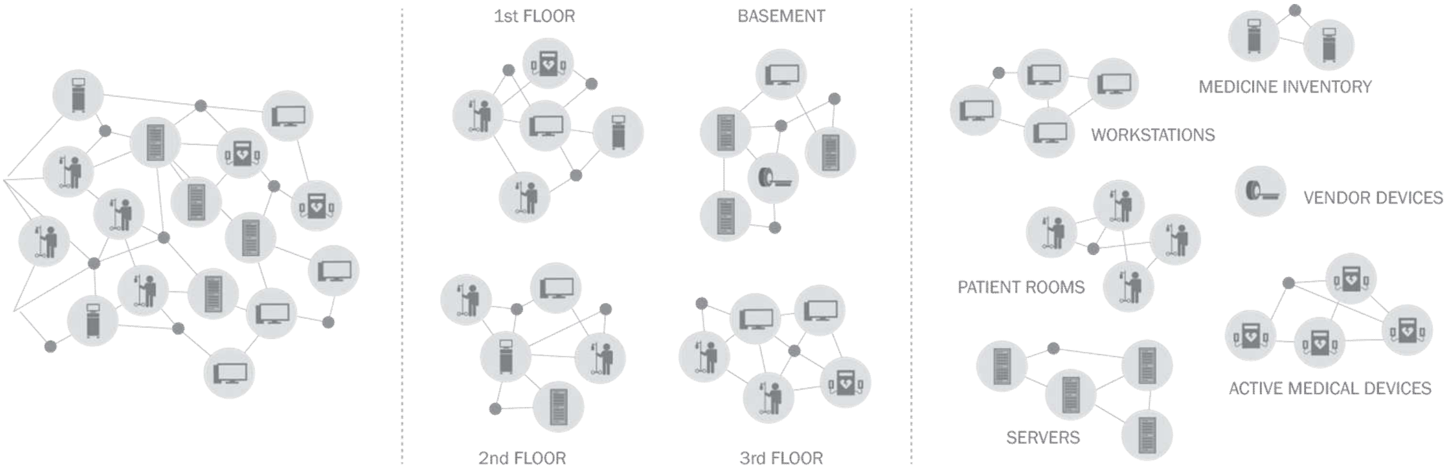

Many networks were flat, with no segmentation.

Other networks were segmented by physical location; poor access control remained the same.

No networks were properly segmented, separating end-systems by purpose, value, or other attributes.

Insufficient/ineffective access controls

Enterprise focus: Technical

A network architecture allows for the effective use of access controls, but they must be implemented. We found that most hospital networks were open between systems. That is, regardless of system type (be in a nurse's station, medicine dispensary, printer, EHR portal, or active medical device connected to a patient), one system could communicate with the other.

It is imperative to take advantage of networking access controls to prevent the infection, spread, and compromise of systems by malicious adversaries, as well as the unwanted egress of stolen data. Networking and security appliances should control the flow of network traffic in as many places as practical, and open networks should not be permitted.

Our blueprint found in this report recommends details for how hospital network access controls should be deployed to create the strongest security posture practical.

Extensive use of legacy systems

Enterprise focus: Business, Technical

A legacy system is one such that it is no longer supported by its creator. Legacy systems could be hardware (in the case of routers, medical devices, and other appliances), operating systems, or software products. Even though legacy systems are not supported, they are often found in use because the system is still operational, costs to upgrade are prohibitive, resource constraints prevent the upgrade process from commencing, or a general lack of personnel or proper management deprioritizes these activities.

The use of legacy systems is dangerous for many reasons. It becomes unlikely that anyone is maintaining the system, that any vulnerabilities found will be mitigated and patch released, or even that any reports of vulnerabilities will be heeded. They therefore exist as *ticking time bombs*.

Our investigations found numerous instances of systems that had reached their end-of-life (EOL). We also found numerous systems that were not EOL, but no longer supported by the expired service contracts for those systems. In a sense, unsupported systems are just as problematic as those that have reached EOL.

Hospitals should ensure that support contracts are up-to-date, that EOL systems are replaced, and that plans are implemented in preparation for such EOL and end of support events that include either upgrading or decommissioning the affected systems.

Weak/unknown controls regarding remote access

Enterprise focus: Policies and Procedures, Vendors

There is an ever increasing movement toward the adoption of remote access technologies so that employees, vendors, and business associates alike can gain access to systems while away from them physically. More physicians enjoy working remotely, or connecting while mobile to the hospitals' systems for which they need access. Likewise, the cost of maintenance is reduced by granting vendors access to the complex machinery and medical devices installed at hospitals around the world. As this trend continues, more and more security vulnerabilities are introduced in to the healthcare ecosystem that must be addressed.

Of the hospitals we investigated, there appeared to be little, and often times no control over parties with remote access to the hospitals' networks. The vendors or remote employees had full control over the machines used to connect, removing them from the hospital IS department's purview, and furthermore, the access granted to these parties was typically far broader than necessary to achieve the mission.

Without control of the remote networks and systems, it is exceptionally problematic (if not impossible) for hospital IS or IT to ensure that those connected systems are safe, and not infected with malware or opening the door for an advanced threat to launch an attack. Hospitals should create strong policies with regard to remote connections, and taking control of the remote devices and/or heavily restricting their access once connected.

Use of custom-built, non-security assessed software

Enterprise focus: Vendors

Several of the hospitals we investigated were found to be using custom-built software, in the form of a web application or other internal system that had not undergone any security assessment. All systems that protect sensitive assets or if compromised could result in the further compromise of higher value assets, must be assessed by a knowledgeable party to ensure that there are as few security vulnerabilities as possible in the software.

Any custom-built software should be built with security as a design consideration from the beginning, and evaluated throughout the process by security professionals knowledgeable with regard to how to circumvent security. If the process has already completed, and software has been deployed, it should still be evaluated after the fact.

Use of vendor provided, non-security assessed software

Enterprise focus: Vendors

Unlike having the capability to assess software built by the hospital in-house, it is not always apparent as to whether a vendor has had adequate security testing of their software or hardware products. It is also not always practical or even permissible for a hospital to conduct this testing themselves. In fact, recent research has shown that security vendors even recommend and in some cases require by contract that users of medical devices deploy them in an insecure manner.[36]

Regardless, it is the hospital's responsibility to ensure the safety and wellbeing of its patients, and thus the onus of ensuring these devices are secure is on them. Since this may not be practical alone, it may be beneficial for hospitals to work together, and with vendors, to have the appropriate security assessments performed and products hardened so that they are appropriate for use in high-risk environments such as hospitals.

Critical uptime issues prevent the implementation/application of security

Enterprise focus: Technical, Policies and Procedures

It is not uncommon for system or software updates to crash, causing critical systems to go down. It is also not uncommon for critical systems to require 100% up-time, meaning that there is simply no time for updates to be installed. However, either of those situations should not justify the postponement or neglect to test or update systems so that their security can be strengthened.

In the hospitals we investigated, we encountered several systems that were not updated due to reasons such as these – that taking the system offline to perform the update was unacceptable, or that application of the update was too risky. Neither of these excuses is acceptable in anything but the shortest term. Hospitals should implement policies that permit the periodic downtime of systems for testing and upgrading. Furthermore, critical systems should have fully redundant environments where, for all intents and purposes, security testing and patch testing can be conducted without interference with the production system. While these steps may seem cumbersome given the duplicative effort compared to the functional deployment, they are essential. No system should go without a security update.

[36] http://www.forbes.com/sites/thomasbrewster/2015/07/10/vulnerable-breasts/#2715e4857a0b6c9f8c8d30e0

Primary attack surfaces on non-restricted subnets

Enterprise focus: Technical

Active medical devices (AMDs), medicine dispensary systems, and equipment aiding in surgical procedures are primary attack surfaces. In several of the hospitals we investigated, some (if not all) of the above were found reachable from the non-restricted portions of the network. That is, many of the same systems used as physician and nurses workstations, printers, and passive medical devices all existed on the same network subnets as these primary attack surfaces.

The fact that these systems are readily reachable creates attack vectors that could allow remote or local, but unprivileged attackers direct access to the attack surfaces most likely to cause patient harm.

Hospitals should create network designs such that these primary attack surface systems may only exist on highly restricted, if not air-gapped, network segments.

Local physical access to critical hospital networks

Enterprise focus: Physical Security

The physical security aspects of the hospitals we investigated were not necessarily in scope, but it was evident that physical security measures were lax, and unlikely to prevent a moderately determined adversary from entering most areas of the hospital. This can be regularly experienced as a patient or guest to the hospital.

The problem arises in this lack of physical security in that patients, guests, or malicious adversaries have physical access to the hospital networks. Most patient rooms have exposed network connections (often for connecting medical devices), and other such ports can be found throughout the hospital. We found that in addition to this physical access, there were often no security measures employed for detecting whether rouge or malicious devices had been connected to the network, meaning an adversary could simply walk up and plug in to the network in order to gain access.

Hospitals should incorporate in to their security design the fact that patients, guests, and potentially malicious parties could obtain physical access to parts of the hospital, possibly connecting to the network from those locations. In addition to addressing this with physical security precautions, there are a variety of digital security measures that can be taken. This should be taken in to account by the facility.

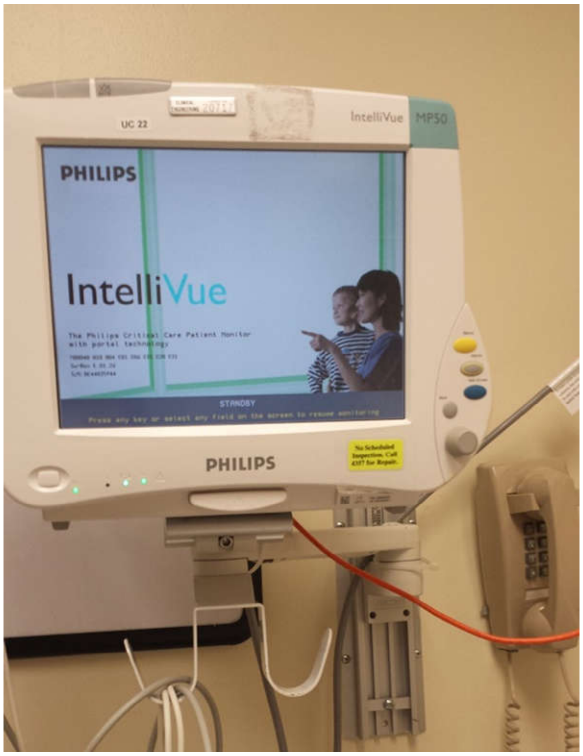

These images were taken from an E.R. patient's room at approx. 5 a.m. We had uninterrupted, private, physical access to the hospital's medical device network for periods of time ranging from 20 minutes to over an hour.

As demonstrated elsewhere in this report, access to these networks can pose life-threatening to patients who are monitored by these devices, or connected to active medical devices also found on these networks.

This sort of physical access to critical hospital networks was common at most hospitals we investigated.

* The Philips IntelliVue system was not part of our research, however, it is a passive medical device of the same category as described in our empirical attack scenario earlier in this report.

Local physical access to systems and devices

Enterprise focus: Physical Security

Our research found that there are many systems within the physical reach of patients and guests at the hospital, including mobile workstations, nurse's stations, unattended terminals, wireless access points, kiosks, and passive and active medical devices. Leveraging physical access to gain privileges on a device is generally a difficult attack to prevent. The fact that physical access is so prevalent within a hospital exacerbates this problem.

With physical access, an adversary could modify or gain control of a device that could later cause harm to a patient, or leverage that access to establish a foothold on the network. Assuming physical connection to the network was not feasible, physically compromising a device that was connected might still be feasible.

Hospitals should include in to their security design the fact that patients, guests, and potentially malicious parties could obtain physical access to systems and devices connected to the hospital network. In addition to addressing this with physical security precautions, there are a variety of digital security measures that can be taken. This should be taken in to account by the facility.

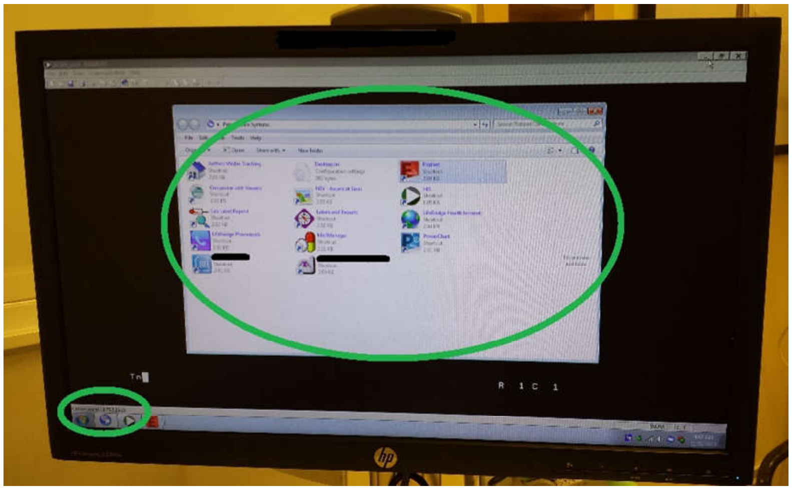

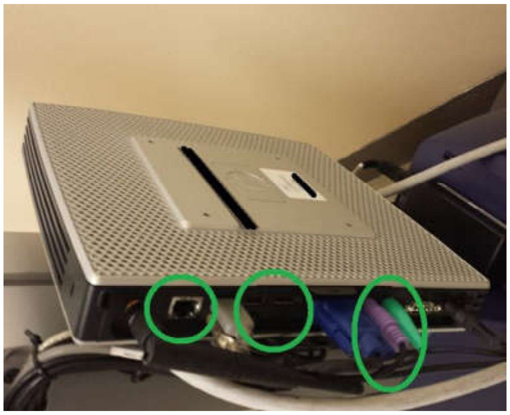

This in-room mobile workstation was left unlocked by hospital staff during an actual patient's stay at the E.R. We were able to keep the session active as long as necessary by interacting with the mouse and keyboard. Many of the applications displayed grant access to EHR, hospital systems, as well as the barcode reader/scanner attached to this workstation. Even though some applications required an additional credential to access, we had full, unfettered access to the workstation itself.

The system hardware was also physically accessible, providing USB access, NIC access, and PS2 keyboard access for which a keystroke logger could be trivially installed.

This sort of physical access to in-room hospital systems was common at most hospitals we investigated. Furthermore, credentials were entered in front of patients, and in some cases a second factor was not required to authenticate.

Credentials entered in the presence of patients/guests

Enterprise focus: Policies and Procedures, Physical Security

Passwords are an important step in authenticating authorized parties to access a system. The workflow at most hospitals, including the hospitals we investigated, required that physicians and nurses enter password credentials in front of patients –essentially exposing those credentials. Repeated password entry then results in repeated exposures, heightening the chance that the credential would be compromised and used maliciously.

Some, but not all, hospitals we investigated used additional forms of authentication beyond passwords, such as a card reader. This second factor is an important step in mitigating access in the event of a compromised password.

When designing the security of mobile systems, physician, and nurse’s workstations where credentials may be entered in the presence of patients, it is important to consider the fact that repeated entry increases the risk associated with a password compromise, and that additional precautions should be taken.

General implementation issues with hospital security

Security implementation issues will be prevalent when there is a poor security design, but beyond that there are very often security implementation issues that are failures to adhere to the intended design.

The following are the variety of security implementation issues we found with the hospitals investigated. We generalize here because the issue categories are more important than the specifics for this report, but in general they illustrate the fact that much has to be done to take hospitals from their current state to a state of comfortable security which can be assumed to mitigate risk.

Use of insecure services

In our investigations it was commonly discovered that vulnerable, known insecure protocols were in use on the network. Or at the very least, they were available for exploit by an adversary. Protocols such as telnet, rlogin, and outdated versions VNC, windows remote desktop, etc., all not only ineffectively protect communications within the hospital network, but leave open doors for adversaries to exploit to gain access to the endpoints themselves.

All organizations should have a written policy prohibiting the use of certain insecure protocols, and when practical include a whitelist of protocols that are permitted. Furthermore, these protocols should be additionally limited and restricted to specific networks. The organization must then regularly audit itself to ensure compliance with this design.

Broken access controls

In our assessments we found that a variety of hospital systems had broken access controls allowing access to parties who should not have access. These came in a variety of forms, from missing authentication fields entirely in certain software, to misconfigured network shares. These vulnerabilities allowed us to gain control over machines we should have not had access, and to bypass the authentication of web applications giving us access to employee and patient records and other sensitive information.

While a written policy helps to avoid misconfigurations, they will still occur in practice. It is imperative that hospitals conduct regular audits of their own systems, and engage outside parties, in order to ensure that these issues are caught and mitigated promptly.

Default configurations

We found that several systems and devices on the hospital networks we investigated had not been fully configured, instead leaving default settings enabled that provided less than optimal security – in some cases permitting default passwords.

These types of issues can be addressed with the proper policies, procedures, and audit, but unless the organization is willing to carry out these tasks, systems will likely remain in their default, vulnerable states. In some cases[37], default configurations are mandated by the vendor, making the proper security choices available to the hospital IS department difficult to interpret.

Shared credentials

During our investigation, it became apparent that a variety of system access credentials were both shared among different staff members, and used to access different systems. This is an improper password usage technique, and in fact was in violation of what password policies did exist for these organizations.

Proper training can best alleviate these issues, but implementation of effective systems for user and account management are also essential in preventing credential misuse.

Unpatched systems

While most of the hospitals we investigated did have patch management policies, in practice we discovered numerous systems without up-to-date patch levels, many of them demonstrating easy to discover security vulnerabilities.

It is most likely a resource constraint issue, as appropriate design was in place that should address these issues. Until hospitals can receive the resources needed, issues like systems remaining unpatched will likely persist.

[37] http://www.forbes.com/sites/thomasbrewster/2015/07/10/vulnerable-breasts/#2715e4857a0b6c9f8c8d30e0

Recommended solutions

One conclusion we draw from our research is that hardly any of these issues can be resolved overnight. The systemic problems will take time to overcome, but this process can be initiated in a positive direction with highly productive and effective results. In some cases, it may take years for a hospital to migrate from its current state to the most secure. Likewise, it will take the industry several years to correct systemic issues and create effective programs for bolstering security on every level, from the device vendor to the hospital to the patient at home.

For the Industry	For Hospitals
Focus on patient health	Follow the blueprint
Create effective regulations	Create a long-term plan
Empower the consumer	Increase funding
Empower the CIO/CISO	Increase security knowledge
Philanthropy	Separate IS from IT

Recommendations for the industry

From an industry perspective, there are numerous steps that can be taken to drive change in healthcare ecosystems toward an overall stronger security mindset. While the final onus of security may appear to fall on one party in particular, this is not necessarily the case. It is the responsibility of all parties involved to participate honestly and strive for the best interests of the end users: patients.

Focus on patient health, not just patient records. For too long the focus has been heavily directed toward patient privacy. There is no question that this is an important factor in protecting patients' interests, however, patient health is the more serious concern and has been overlooked far too long. It is necessary that the industry as a whole adopt the stance that patient health is a digital security concern as important as (or more important than) patient privacy. Until this industry-wide change is made, it is unlikely that the risk to patients will reduce.

Avoid (or create effective) regulations. If government regulations are necessary, they must have teeth. For almost two decades, HIPAA has been ineffective at protecting patient privacy. HIPAA has, however, created a system of confusion, fear, and busy work that has cost the industry billions of dollars. The result is excessive distraction and expense that could otherwise be put toward proper security of patient privacy and health; the true end goal. Punitive measures for compliance failures should be treated carefully, so as to incentivize the search for non-compliance issues to fix, and not dissuade the search in the first place. Healthcare organizations should be rewarded for pro-active security work that protects patient health and records.

Empower the consumer (comparative ranking). Consumers have little knowledge of where their information goes, let alone whether or not those information custodians practice good security. Likewise, they do not have the requisite skill or resources to determine if receiving care at one facility versus another affords them a greater level of protection. An industry-wide comparative ranking system —if done right— would empower the consumer to make informed decisions about the security of their health and privacy when choosing a provider.

(Re)Empower the CIO/CISO and other executives. Much as consumers have little insight or control over the care providers they choose, decision makers at hospitals have little insight or control over the security practices of their vendors. These decision makers are left with choosing between vendors boasting the better buzzwords, but have no real knowledge of any truth behind the assertions made, or even if they're the appropriate assertions. The result of a bad choice could be a permanent smear on their facility. Instead, decision makers should be empowered to make informed decisions about the products they procure, and the vendors they work with. Third-party security assessments by experienced professionals can lend to this empowerment, if vendors are required to produce such evidence. Decision makers can then accurately use security as a decision point when choosing vendors.

Philanthropy. Good security can often be cost prohibitive. Much like an endowment, grant, or donation of funds that could be used for medical equipment or staffing, these funds can appropriated to elevate the security posture of an organization. For instance, a hospital grant could allow for seasoned security professionals to be employed by the hospital, or for an organization to vet the security of medical devices used by a multitude of hospitals.

Recommendations for hospitals

For healthcare facilities –the target of this research– there is incredibly important work to be done. These are the organizations directly responsible for patient health, and if ever patient privacy was *not* a priority in the mission of a healthcare facility, there is no question that the security of patient health *is* a priority. It is the responsibility of the facility to protect patients, and these are some recommendations on how to better do so.

Follow the blueprint. In this report we've included a blueprint for better healthcare facility security. The personnel at these facilities responsible for security should adopt this blueprint and begin to act on it. Our blueprint is adapted from other common standards, but most applicable for healthcare facilities and taking these facilities from their current insecure posture to one that is much more secure.

Create a long-term plan. The executive and board levels must recognize that security is not an overnight patch-and-fix activity. Appropriate security always involves long term planning, much like any other business unit, including information technology with which it is most commonly associated. While there are immediate needs to patch and plug holes, efforts such as increasing budget, restructuring the organization, rearchitecting the network, or hiring additional personnel, are all long-term efforts. These long term plans should be created as soon as possible, and updated as the technology, threat, and defense landscapes evolve.

Increase funding. The blueprint found in this report describes what is needed for a healthcare facility to achieve strong security. The impeding factor in obtaining the majority of facilities is a lack of funding. Nearly all aspects of the blueprint, from needed equipment, to needed personnel, to needed training, to needed

Compliance Threatening Security

During the course of our research, multiple hospitals were concerned that we would uncover a security breach for which they would then need to report; the result of which could be significant fines among other penalties. In fact, these facilities indicated a reluctance to engage an outside security firm altogether, for fear that if as a result of those engagements a breach were found, that they would suffer massive penalties.

While true that non-compliance with regulation must have repercussions, lest those statutes have no teeth, they become fantastically ineffective (and even counterproductive) when the punitive measures dictated by those statutes are applied to hospitals making good faith efforts to enhance their security.

Regulatory statutes should **reward** healthcare facilities for taking proactive measures to enhance security; not penalize them during the process. Otherwise, as we have seen, these facilities will prefer to *do nothing* over taking the much needed steps to protect their patients' health and privacy.

Our recommendation is that any statute that commands penalties to be levied on an institution for non-compliance or confirmation of a security breach, waive those penalties if the violation or breach is uncovered during a proactive security assessment.

consultation require funds. In conjunction with a long-term plan, funding can be increased and the most appropriate security activities prioritized ahead of time. There are minimums by which security will be entirely ineffective unless they're met, but these will differ between facilities.

Increase security knowledge. The facility should continuously seek to increase its *security knowledge* through regular training and by augmenting their team with security professionals, either through hiring or engaging outside consultants. This increase in subject matter knowledge is crucial for an organization to competently design and execute a security strategy.

Separate Information Security from Information Technology. While both areas involve technology, it is inappropriate to treat information security as an information technology effort. The two departments are forever in conflict, and those conflicts must have an appropriate path of escalation to a decision maker with responsibility, competency, and knowledge in both areas.

Part III: Healthcare Facility Security Blueprint

For most healthcare facilities, it is not a question of *am I secure*, or *how secure am I*, but of *how do I get there*? This question of how to get from where they are to a point of security readiness is difficult, and the further that distance the more daunting this task becomes. When the task at hand is discouraging, it is prone to delay, waste, and failure. Our hope is that this blueprint can offer guidance (and comfort) for healthcare facility officers in charge of information security and the patient's health.

Timeline

Varies/Perpetual. There is no comprehensible way that abstract timelines can be placed on real world facility initiatives. All facilities are different, as are staff, budgets, skill set, and so on. The two most is important things to realize regarding timeline are that security takes time, quite possibly a lot, and that the activity is recurring and perpetual. The process itself to reaching security readiness could take anywhere from several months to several years. Proper planning ensures that the highest priority items are addressed early (reducing as much risk early on as possible), and that the plan can be revised as changes are made in the future. When readiness is then reached, the facility must realize that the problem is not solved, or that work or expenses will stop. Security is an ongoing process; readiness is reached when the recurring processes for planning, training, system hardening, audit, updating, monitoring, etc. are underway and effective.

Cost

Varies/Recurring. As with timeline, cost is far too difficult to estimate usefully while remaining generic. This depends on the facility's current deployment, staffing, threat model, timeline constraints, etc. The cost also depends on what the facility deems a security related cost. Sometimes staff is considered to fall under one budget while equipment and training fall under others. Each could be attributed to security, but it would be misleading to speculate as to total budgets. In our blueprint below, we explain what is needed, and the general cost of those initiatives. This can aid a facility in its planning and budgeting.

Process

We've broken the process of migrating from where a facility is to security readiness in to the following phases. These are listed in the ideal and most intuitive order, but in practice healthcare facilities will need to enact portions of each of these categories simultaneously and out of order. This is why proper planning at the beginning and throughout is so important. While it may be unavoidable, the further out in the order of phases that the organization seeks to act on without first acting on the previous steps, the more likely there will be waste in duplicative or lost effort. For instance, hardening and auditing systems prior to having a defined network policy may result in those systems needing to be reconfigured and audited again, though it is reasonable to expect that the hardening of critical systems would take precedence in a long-term security plan, even before a well-defined policy was prepared. It is the nature of this process.

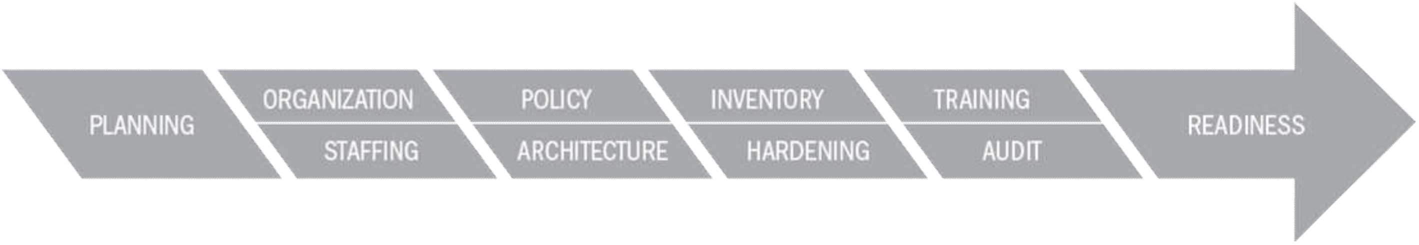

Planning

Planning is the intuitive and obvious first step to this process. Whether it involves a full blown long-term plan, or an immediate action plan to plug the most glaring holes, a plan is required. It is costly and wasteful to act on security arbitrarily.

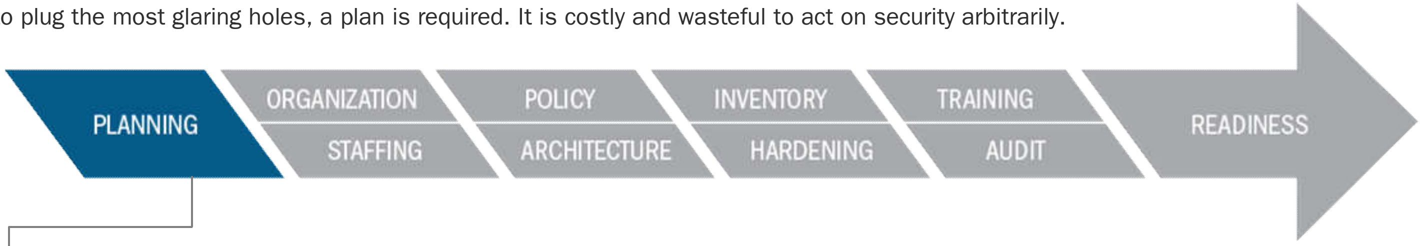

Understand the blueprint

The facility should use this blueprint as a guide to creating long-term and short-term plans. It is not a one-size-fits-all blueprint, but it should get you a majority of the way there with your organization prepared to fill in the gaps.

Create a threat model

The facility plan should include a threat model. This should identify the assets under the control of the facility, including primary (patients, patient records), secondary (samples, certain records, some servers), and tertiary assets (other systems in the ecosystem). The model should identify all of the relevant attack surfaces, including primary (medicine dispensaries, active medical devices), secondary (work orders, test results, servers), and tertiary attack surfaces (mobile devices, networking equipment, email). The threat model not only helps in planning, but coordinates the vision of staff members, and provides for the lower-cost participation of outside consultants. It helps align parties with the mission.

Understand the risks

The facility should also spend time evaluating the highest value targets, whether they are patients, certain kinds of records, or even certain collections of records. In some cases, a research facility might also have valuable assets in the form of biological samples that need to be preserved, or the integrity of a study. Also, using the threat model the facility should identify the highest risk attack surfaces. These are likely the devices that are closest to the assets, like active medical devices, but also can include certain servers or networking equipment that could play a role in an attack. All of this must be mapped out for the organization in order to plan, budget, and prioritize actions.

(Re)design security framework

The threat model and risk evaluation will hopefully shed light on the gaps in the current security design that must be filled in order to achieve stronger security. The facility should consider all aspects of security, not just the traditional perimeter, physical, and people attack surfaces, but also the variety of attacks that originate from within.

Perform gap analysis

With a security framework designed, it is then possible to look at the existing network and perform a gap analysis to determine what needs to change in order to have adopted the new design. These are more or less the steps necessary for the plan to succeed.

Create long term plan

With each of the above completed, the facility can create a long-term plan with actionable short-term goals. The plan should prioritize the security of primary assets, and incorporate current and future budgeting requirements, and at a minimum address all of the categories below.

Organization

The healthcare organization, in all likelihood, will need to be reorganized with new processes before an appropriate security posture can be established. This can be a difficult hurdle to overcome for many facilities, but it is nevertheless essential to the process.

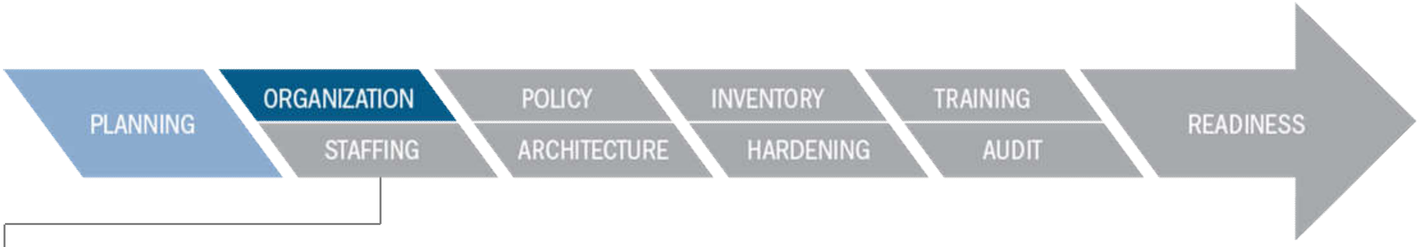

Information Security (IS) separate from Information Technology (IT)

In most healthcare facilities, information security duties fall under information technology, but this is an inappropriate structure. The directives for IT and IS are in conflict, and when these missions fall under the same department, the decision making responsibility often gets pushed to the lowest level, and the criticality of those decisions overlooked. A more appropriate structure is for IS to report to a CISO, CIO, or CEO directly, rather than then head of an IT department.

Biomedical engineering (and other technicians) under IT and/or IS

In most healthcare facilities, the biomedical engineering department (BioMed) is distinct from IT and free to operate within the hospital at its discretion. Likewise, it is common for departments with specialized equipment (e.g., Radiology) to have technicians with unfettered access to the hospital network to make changes or connect devices. The fact is, these activities involve the facility network, technology systems, and networked devices that interface directly with patients. Without oversight by IT or IS, it is unlikely that any real security can be built around these departments. Instead, IT, IS, or both depending on the facilities processes, should have full knowledge and oversight of any activity which modifies or connects equipment to the facility network. Furthermore, IS is responsible for the health of patients as it relates to digital attacks, and thus is responsible for implementing, auditing, and monitoring the security processes put in place for BioMed or these other departments.

The following is a recommended, generic organizational chart for a hospital.

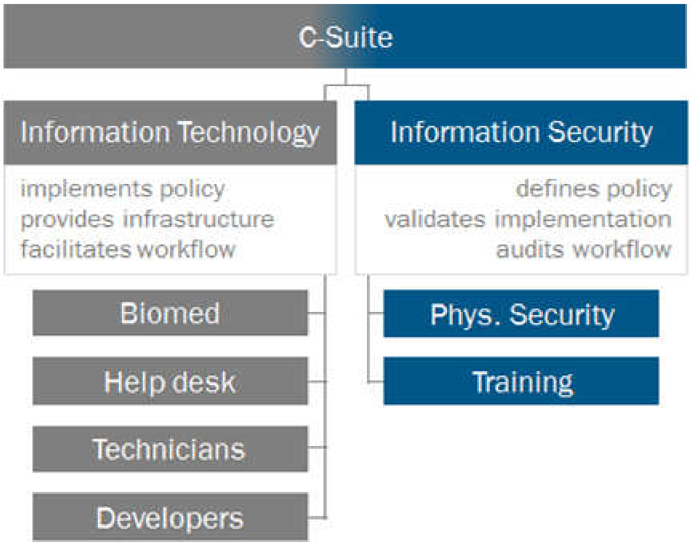

Staffing

No matter what, achieving the security goals of the healthcare facility and following through with this blueprint requires appropriately trained people to carry out the tasks. Until people resources are available, little material improvement will be seen. Exactly how many people will certainly vary between organizations, and will depend on facility size, network size, desired security posture and speed at which it is obtained, and whether or not the facility intends to conduct its own training, auditing, or other aspects of this program that are sometimes outsourced.

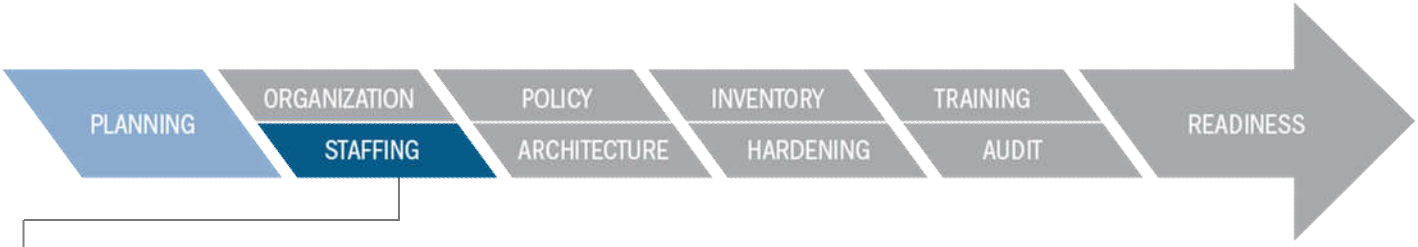

Roles

Staff roles should be differentiated between management and design roles, and implementation of security roles. This does not mean that individuals cannot participate in both roles, in fact, they should. However, the two role-sets require different experience and skill levels, and they scale differently as security is put in to practice and as the organization grows. Management and design roles require greater experience, typically 2-5 years of relevant industry experience with an emphasis on soft skills, while implementation roles *can* require less experience if properly managed, but are highly technical. As more security is implemented or the organization grows, the need for implementation-specific roles increases at a higher rate than management-specific roles, and the facility should plan for this accordingly. However, it should be noted that the security mission will not be successful without both, and that the management role is likely required earlier in the plan.

Full time vs. part time

The facility may leverage personnel across other roles within the organization, or outsource tasks to third parties, but it is imperative that a core group of employees (even if just one) be responsible solely for these security initiatives. For instance, IT managers may have security implementation roles, and HR managers may have security training roles, but there must also be security-specific personnel. It is far less effective to have all security responsibilities delegated to individuals with only a part-time responsibility for security. It is far more effective to have a core team of security-specific roles with the ability to delegate and oversee that other personnel with part-time security responsibilities carry these tasks out correctly.

Quantity

Each facility will have to determine that which is the maximum valuable staff before diminished returns or waste is experienced, but in our experience, hardly any facility is close to reaching this situation. There are, however, effective minimums that we'll discuss here.

Management Personnel	Small (<20 beds)	Medium (21-50 beds)	Medium-Large (51-99 beds)	Large (100-299)	Very Large (300-499)	Super Large (501+)
Minimum	0.5	0.5	1	1	2	2
Recommended	0.5	1	1	1	2	3

Implementation Personnel	Small (<20 beds)	Medium (21-50 beds)	Medium-Large (51-99 beds)	Large (100-299)	Very Large (300-499)	Super Large (501+)
Minimum	0.5	0.5	1	1	2	3
Recommended	0.5	1	1	3	4	6

The above tables show that there are stricter minimums for smaller facilities —often specifying at least one individual per role no matter what— but the recommendation above the minimum is to target roughly 1 information security officers per 75 beds. Since facility metrics may vary, the following table can help identify a staff target in atypical situations. Keep in mind, these recommendations could vary greatly in practice depending on the structure and efforts of the organization, and may not be appropriate minimums for some facilities like research facilities that may have higher staff-to-bed count ratios.

Personnel	Beds	Total Staff	IT Staff	End Systems	$ Revenue	Patients
1 ISO per	75	800	8	400	$300m	4,000

Policy

The healthcare facility's policy is that which formally defines security, specifying goals, requirements, and procedures throughout the organization. These are also the information sources by which assessment and audit can be conducted. Without a defined policy, implementation of security is arbitrary. The following diagram shows how the components interrelate.

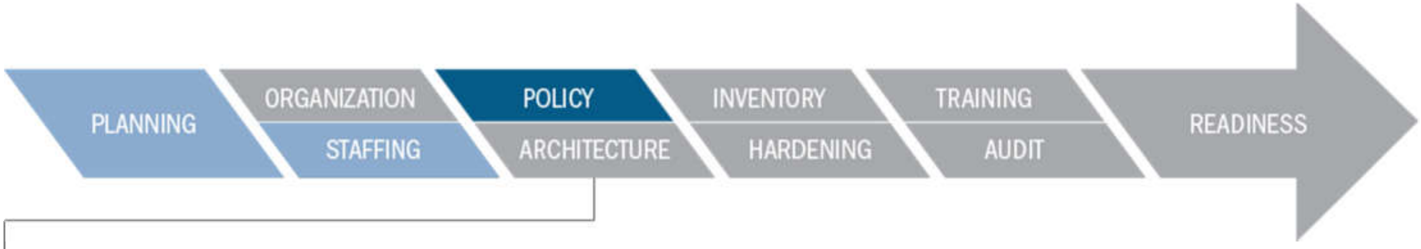

Policy (design) is the structure of how security is meant to be. Implementation (deployment) is the act of putting the prescribed policy in to effect. Assessment is the determination as to whether or not the policy is good, or the implementation is effective at upholding the policy. Audit is the act of verifying that the implementation is true to the policy.

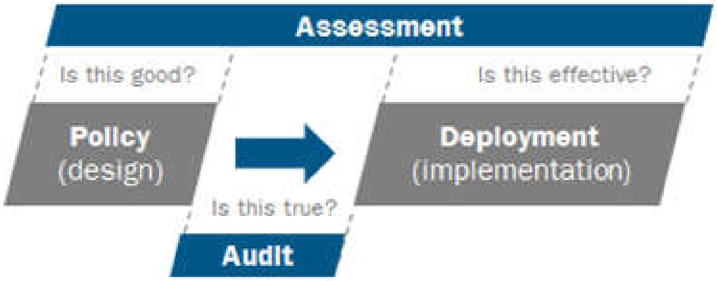

Formally define security

By formally defining security, it aligns all parties in the security process and allows anyone reviewing the policy to ensure complete understanding of the mission. The formal definition should incorporate the previously developed threat model, security goals, those parties involved in the process, etc.

Needed policies

The policy will have many subparts, and should define as many of them in as much detail as possible. Ideally, the policies of the facility touch on all areas and embody the formal definition laid out previously. The following are typical policies that would be appropriate to include in a healthcare facility security policy. Each policy should define what is and is not permitted, both at a high level and down to specifics. It is important to articulate the purpose of the policy (high-level) in addition to the specifics, so that interpretation of the policy can be best aligned with the policy creators.

Recommended Policies		
• Physical security policy • Patient access policies • Guest access policies • Network security policy • Personnel security policy • System users / account management • Software security policy	• Remote access policy • Appropriate usage policies • Personally owned devices policy • Security training policies • Security reporting policy • Email/web policies • Web browsing policy	• Medical device policy • Sample collection policy • Test collection policy • EHR handling policies • Workflow policies • Endpoint security policies • Information logging policies

This is not a complete list of policies, but should shed a decent light on the types of policies that a healthcare facility should create. The facility should first create a list of all required policies, and then begin defining them.

Needed procedures

Undoubtedly, the policies will indicate a variety of specific procedures that must be carried out when different events occur. By defining the procedures, it aligns all parties with regard to actions, expectations, the standardization of deliverables, etc. The following are typical procedures that are appropriate to include in a healthcare facility security policy. Each procedure should specify what is the procedure's purpose, under what conditions the procedures should be enacted, who should be contacted and involved, and the step-by-step process that should ensue.

Recommended Procedures	
• Employee termination. • The addition and removal of systems on the network. • The updating of medical devices. • The installation and updating of software.	• The reporting of security incidents. • How audits should be conducted. • How new employees are trained.

This is not a complete list of procedures, but should shed a decent light on the types of procedures that a healthcare facility should create. The policy should indicate which procedures are needed, and the facility should then create them.

Audit prep

The policy dictates the implementation; audit is the process by which the implementation is verified to uphold the policy. The policies themselves should lend themselves to auditability, and the facility should create an audit procedure. The audit procedure should indicate which areas of the policy are not auditable, which are auditable, and in what manner the audit should take place. Additionally, the policy should describe the frequency of audits, the management and dissemination of results, the procedure for correcting noncompliance issues, how follow up to ensure compliance is conducted, and other actions that might result from this process.

Review

Policies are not write-once, implement forever –they are subject to change. Policies should be reviewed at least once per year. They should be adapted to include new and changing technologies, changes in the threat model or goals of the facility, changes in available resources, etc. This review process should be part of the policies as well.

Architecture

Much like you would design a building to be both functional and secure (constructed so that access controls and guards can be placed at various chokepoints), the network too should be architected with security in mind. While in practice, a digital network is not like a physical building, the network is for all intents and purposes the digital structure by which all information is commuted. The notions that it should be functional, and that security should be built in and scalable over time are essential to a strong network design. Healthcare facilities come in all shapes and sizes, and by virtue of their workflows, mission, capabilities, and specialty, the digital networks that support them will be of all sorts as well. What is important is to architect the network with security in mind while planning workflows, network segmentation, the purchase of networking and security appliances, how the network is administered, and how remote and physical access to the network are handled.

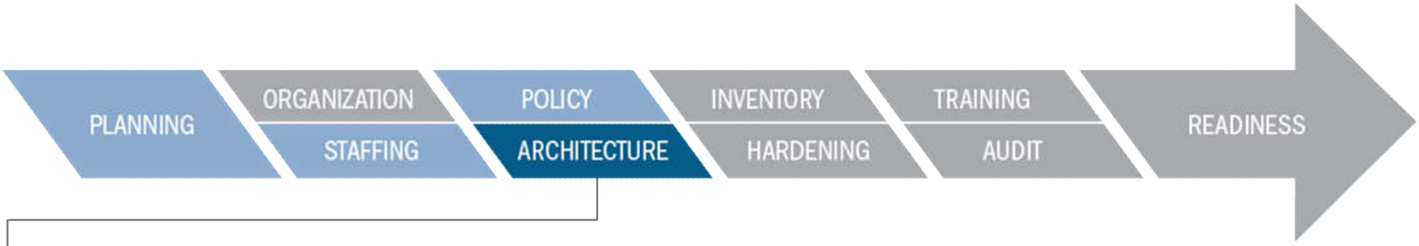

Information workflows

The flow of information through the hospital, if pertinent to the security of patients or patient records, must be documented, including all systems involved in the workflow, the network paths they traverse, and the networking appliances they flow through. Once this information is known, the network can be designed such that workflows with similar security constraints can benefit from the same security policies and administration, while workflows that have different security requirements, or introduce additional risk to the network can be separated. Remember to keep things flexible as workflows will change. When they do, it may be required to rearchitect some of the network.

Segmentation

Proper segmentation or zoning of the network is essential for a strong security posture. Segmentation allows the separation of networks based on security requirements (often driven by business purpose or asset value), but also helps enforce workflows, control the ingress and egress of information, provide vantage points by which monitoring and other defense-in-depth measures can be deployed, and prevent the spread of malware or the progression of an attack in progress.

It is common and likely that the healthcare facility has already implemented an inappropriate network architecture. Rearchitecting and then redeploying a large network can take a long time. Nevertheless, it is necessary. Long-term plans should account for the incremental migration from one network type to the new segmented network.

Examples of BAD network architectures

- Open access architecture (i.e., most machines can talk to most other machines unhindered).
- Geographically based architecture (e.g., zoned for building 1, building 2, or floor 1, floor 2).
- Technology specific architecture (e.g., zoned for all workstations, all servers, etc.).
- Appliance specific architecture (e.g., zoned for different wireless LANs vs. wired LANs).

Examples of good network architectures

- Asset-based segmentation (i.e., different classes of assets reside on different zones).
- Purposed-based segmentation (e.g., radiology, biomed, nurse's stations, administration, etc.).

Network appliances

As part of determining the network architecture, the choice of network appliances should be made to support the architecture, and not the architecture determined to be supported by the existing appliances. The choice of appliances should consider the security features of those products, such as logging, layer-2 filtering, access control lists and enforcement; the supported lifetime of the appliances, so as to avoid having EOL or un-upgradable devices; and the budgetary effect of these devices as it affects security. Often times networking appliances are essential to the IT infrastructure in general. By selecting devices with strong security controls, it eliminates the additional budgetary needs for appliances with these features elsewhere.

Security appliances

In addition to networking appliances, there is the need for additional appliances to aid in the application of security to the network. These appliances include firewalls, intrusion detection systems, log aggregators, network scanners, etc. For the most part, these are not needed for the functional operation of the network, but are necessary for maintaining strong security. They are likely to be accounted for in of a separate budget, unlike networking appliances with security features that may be considered direct IT expenses, but even so, some log aggregation systems are both useful for IT and IS purposes.

Administration

The healthcare facility should establish an administration protocol by which designated administrators are responsible for the administration of security and other technology under the guidelines of the information security department. Proper and secure administration is necessary to uphold the integrity of the network, and thus the policies and procedures for this must be effective. These procedures should be carried about by designated personnel on designated networks whenever possible.

Remote access

The facility should create a restrictive and controlled remote access policy to round out the network infrastructure. It is common for remote employees and vendor technicians to request (and sometimes require) remote access to facility systems in order to do their jobs. If remote access is permitted, the facility should decide under what restrictions that access may be allowed. E.g., whether or not the facility requires purview over the connected devices, or if connected devices are highly restricted once connected, or if access is simply denied. Everything from the policy, to the technology, to the upkeep and maintenance of account management are all essential when permitting remote access from and over hostile networks.

Physical access

Beyond the dimensions of the digital network is the fact that parties in the physical world may have physical access to the network, networking components, or endpoints on those networks. The available controls for mitigating physical attacks are limited, but the facility's security policy should nevertheless address these issues with the best security approach possible, and document and understand the risks associated with the security limitations. It is common for patients and guests (and therefore malicious parties) to have physical access to hospital rooms and the equipment within them. Networks should be architected to account for these threats.

The following diagram shows what a typical facility may do depending on specific needs. This architecture is not necessarily appropriate uniformly across all healthcare facilities, but it is a strong starting point for what is typical.

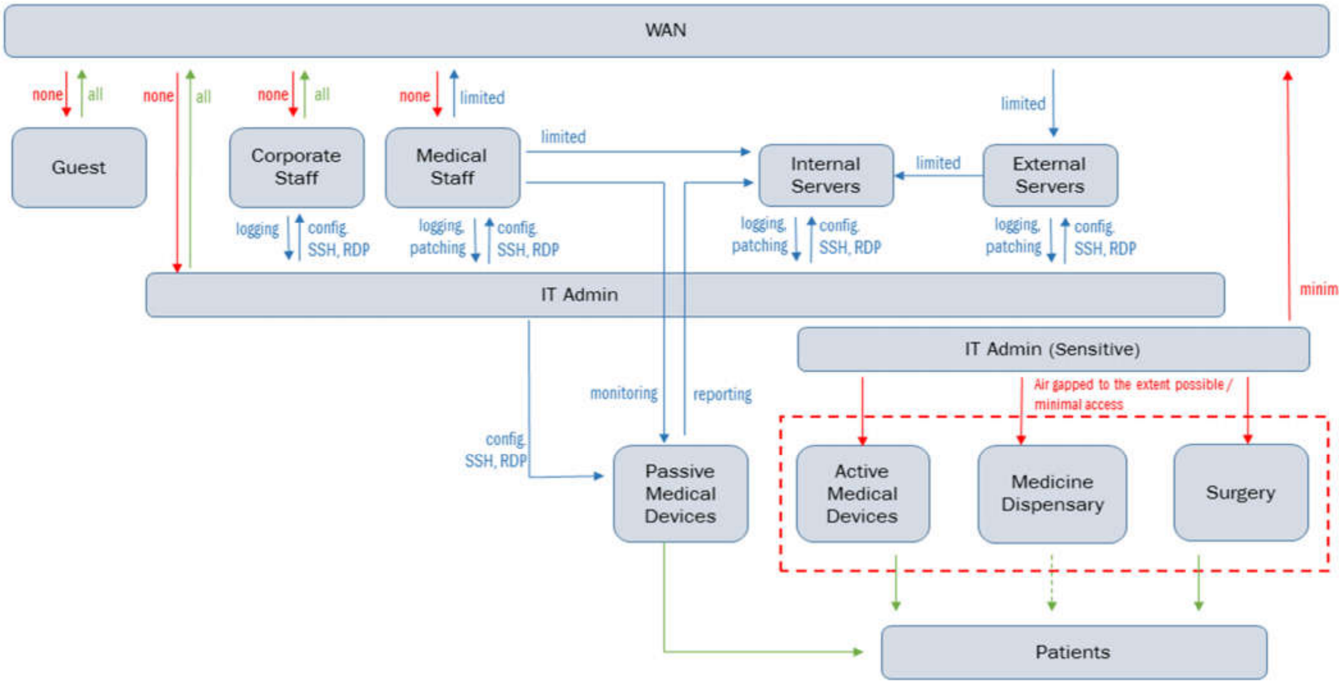

In the figure above, notice the following distinctions:

- A guest network is provided for patients, visitors, staff personal devices, and other systems. These devices should never be connected to the facility networks except in this fully segregated way. If the facility supports information kiosks, retail stores, or other non-facility entities, though they may not be guests in the same untrusted sense, they should be treated as untrusted and segregated in this way, off of the primary facility network.
- Because medical staff deal with different information than corporate staff (e.g., patient medical information vs. accounting and human resources records), these two types of systems should not coincide within the same zones. They should be separated, with different access controls, policies, and monitoring implemented on each.
- The IT administration network zones should be entirely separate from corporate and medical staff. These zones should contain administrative workstations, log aggregators, directory services, patch management systems, etc. All network appliance configurations should take place from within these zones.
- Zones that contain appliances that could directly affect the patient must be highly restricted, controlled, and monitored. These are zones that include active medical devices, medicine dispensary systems, and any equipment necessary to conduct surgery. They should be separated from each other, and air gapped or permitting only the bare minimum services in order for the facility to operate. Since these zones are the most sensitive, they should be administered from a separate, higher security administration network zone, which has additional limitations on how it can be reached.

In practice, there will be many more security decisions to be made than the simple diagram above. For instance, the above diagram does not address remote access network zones. Should these be required, it would be wise to implement them alongside of these existing zones, but in a manner that provides for stricter monitoring and access controls, while limiting access to other systems of its class.

Inventory

Awareness of the digital infrastructure is necessary in order to continuously provide strong security. It is important to always know the types and specific users, appliances, endpoints, and software in operation on the network, as well as how they are connected and communicating. This knowledge helps confirm the security plan is being practiced, helps with the investigation of potentially malicious events, and allows for security procedures to be carried out quickly and effectively. When building an inventory of the infrastructure the following areas must be considered:

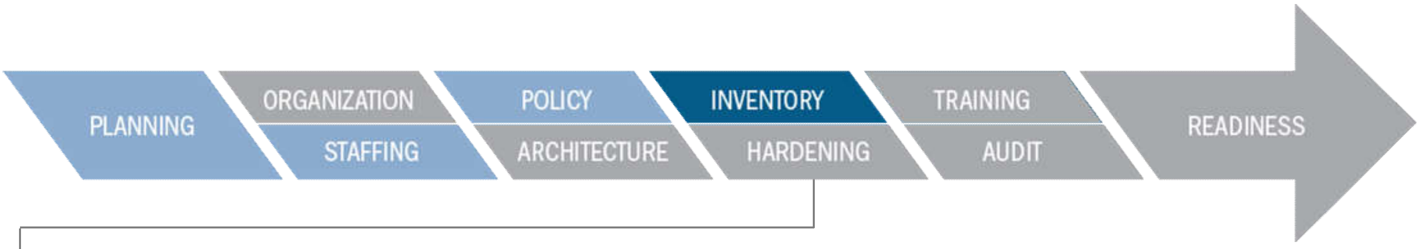

Employees

Any person with access to the digital network should be part of a list, database, or inventory that records the important information about that employee, but also the security-related access they have to the infrastructure. Such information may include system credentials, physical access cards, printed documentation, associated licenses, etc. These inventories help execute security procedures without worry, such as employee termination, changes in employee roles, and upkeep of employee training.

Vendors

Vendors are not all that different from employees with regard to the digital access they may have to the infrastructure, but they have additional information as well. In addition to the credentials and system access that should be documented, vendors may own certain equipment, have service contracts with the facility that may expire, or may have exceedingly limited access. Keeping track of all vendors supports the security initiatives of the organization.

Systems

The facility should inventory any system with an IP address, including endpoint workstations, mobile devices, networking appliances, servers, printers, and medical devices. Inventories should contain system information, patch level, designated administrator, and the functional use and asset information related to the system. This information is important when restructuring the infrastructure, decommissioning systems, ensuring the latest patches have been applied, performing audit, or performing incident response.

Software

Software, much like systems, must also be inventoried. In order to prevent malicious compromise through a system's software, it is essential to make sure that all software within the infrastructure is documented, including the software vendor, patch level, and license information. This allows for control over endpoints with regard to installed software, facilitates the audit of all software, and ensures that software has not reached an EOL or end of support situation.

Devices

Similar to systems with an IP address, all other devices should be inventoried as well. In an ideal world, every piece of equipment from keyboards to USB sticks would be inventoried, however, this is difficult in practice. At the very least, the facility should inventory all devices that connect to a digital system reachable from the network and that affect workflow or could affect patient health or records. For example, active and passive medical devices, barcode scanners, automated environmental controls, scales, etc.

Hardening

Systems and devices deployed on the network must be hardened, i.e., configured to be in the most secure state that still supports the needs of the infrastructure. Software, firmware, threats, defenses, and the healthcare facility's workflows will change. At any of these times, it is usually necessary to revisit the configuration of the network and its endpoints to ensure that they are in the most secure state possible.

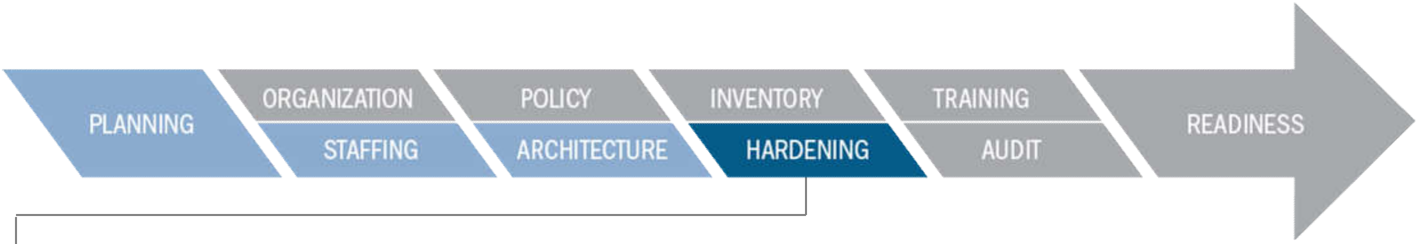

Configuration management

The facility should create a process for managing the configurations of all systems. This process should convert the relevant security policies in to actionable configuration guides for systems on the network, and then implement those configurations. The facility should keep records of all systems that have been configured according to the prescribed guidelines, and if ever these configurations should change, the changes should be documented.

Endpoint security

The facility should create a method for securing the variety of endpoints found throughout the infrastructure. This defense in depth measure could include anti-malware software, local firewalls, group policy settings, or even augmenting systems with additional second-factor authentication tools.

Monitoring

Maintaining a strong security posture requires constant observation in order to identify weaknesses, potential attacks in progress, or attacks that have been successful in order to quickly eradicate them. The facility should have a robust logging, monitoring, and response procedures in place and practiced.

Training

Training is not just the teaching of a new skill, but the creation of awareness, introduction to ideas, and repetitive instruction required to instill a new behavior. Security training should be administered company-wide, with variations between program materials depending on the recipient, and should be on-going. It is not only important to receive training, but to receiving it on a recurring basis.

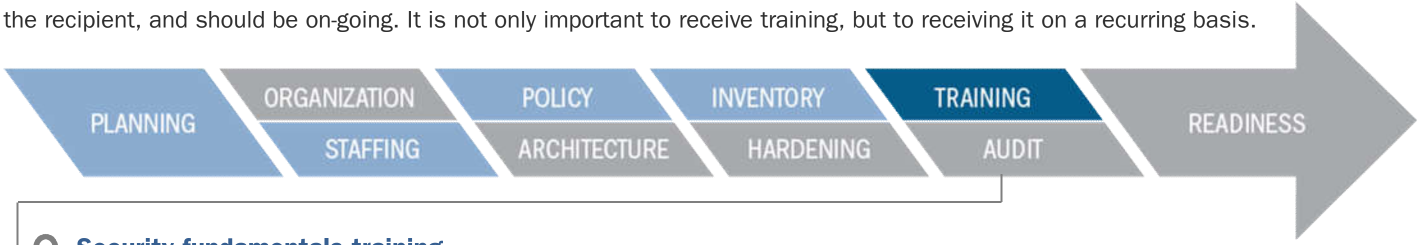

Security fundamentals training

All facility employees should undergo some level of awareness, or security fundamentals training. Similar to the company-wide need for an understanding of ethics, workplace sensitivity, or patient privacy, understanding the basics of how security works and its importance goes a long way in rounding out the corporate security posture. This breathes security understanding into the employee-base from the ground up, and the organization can expect fewer mistakes, and greater diligence in the reporting of security incidents.

All staff should receive a minimum of 4 hours (recommended 8 hours) per year of security fundamentals training. Members of the information technology department, technicians, and other staff regularly administering technology should receive an additional minimum of 4 hours (recommended 8 hours) per year of security fundamentals training specific to information technology.

Policy training

All facility employees should undergo security policy training specific to their job role. This ensures that each employee knows what is prohibited and required on the network, as well as the procedures relevant to their function. There will be great overlap with topics describing how to report suspicious activity or overt security events, and specific procedures can be taught more exclusively, such as handling firewall configuration changes or installing software updates.

All staff should receive a minimum of 4 hours (recommended 8 hours) per year of security policy training. New hires should receive this training within 3 months.

Business training

The decisions makers at the healthcare facility, including C-level personnel and the board, should undergo training regarding security from a business perspective. This training should inform decision makers of the actual importance of security, how it is designed and implemented, and how it is maintained, dispelling myths and helping to better grasp the cost/risk tradeoffs in working with security.

All decision makers should receive a minimum of 4 hours (recommended 8 hours) per year of business-level security training.

Subject matter training

The information security personnel, and possibly even information technology personnel should undergo recurring training to keep them abreast of current attack and defense trends and techniques. This training is more detailed and specific to security theory and implementation.

All security officers should receive a minimum of 16 hours (recommended 40 hours) per year of security subject matter training. It is also recommended that information technology personnel, technicians, and other staff administering technology to receive at least 16 hours of subject matter training per year.

Assessment and audit

Assessment and audit are often confused, but both are necessary to ensure a strong security posture. Assessment is the determination if something is *good* or *effective*, while audit is the verification that something is true. Semantics aside, all healthcare facilities must engage in both activities on a regular, recurring basis.

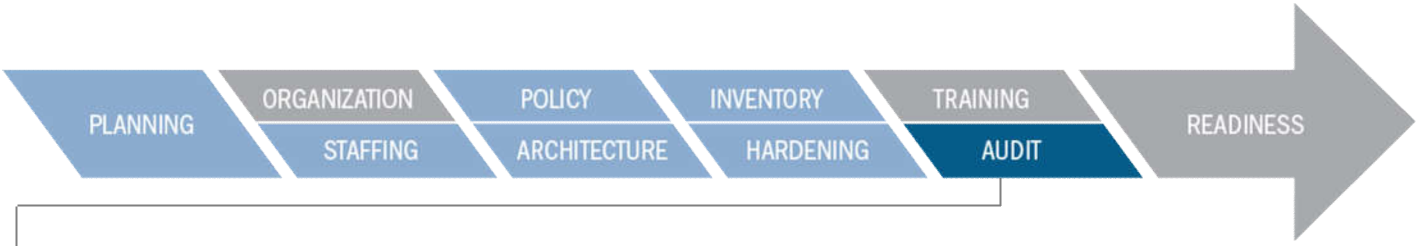

Internal assessment

The healthcare facility security personnel should assess the current state of their own infrastructure, including the plans, threat model, design, and implementation choices on a regular basis. This should happen no less frequently than once per year.

Outside assessment

The facility should also seek to engage outside consultants to vet the chosen plan, polices, design, and implementation. This activity eliminates corporate politics and other ingrained, potentially bad, security opinions within an organization. These engagements are best when ongoing and should happen in full no less than once per year.

Internal audit

The policies and procedures of the healthcare facility should have numerous audit activities, such as the pruning of inventory lists, verification that defunct accounts are indeed removed, checking to ensure systems are patched, and that networking appliances are configured in accordance with the policy, etc. These types of audits should be performed regularly, some occurring daily, monthly, quarterly, or of longer terms.

Readiness

Once policies, procedures, training, and implementation are in place, the healthcare facility is prepared to practice readiness activities. These activities are akin to disaster recovery plans, and include incident response plans, red teaming, and other contingency plans. This is the exercise portion of the security process that the facility must enact, as it ensures timely, effective, and less costly responses when certain events occur.

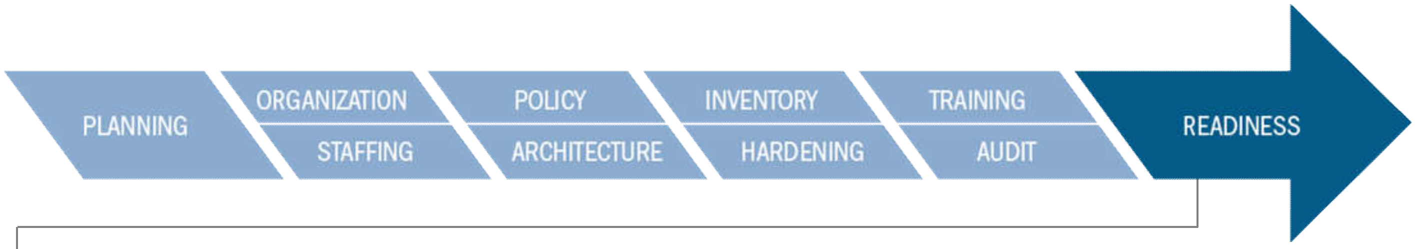

Incident Response

Security breaches will occur, and it is important that the facility staff is prepared to deal with these situations. The procedure0s for the organization should include what to do in the event of a security breach, and all of the relevant details and research having already been completed by that time: such as, who to contact in the event of a breach. Incidents are likely to occur frequently enough that response to those incidents will be naturally practiced, but uncommon breaches should be simulated as well, in order to prepare staff.

Disaster recovery

Security should be built in to any existing facility disaster recovery plan. These plans should be reviewed and assessed as part of an assessment procedure. While disasters are unlikely, they should be simulated so that the security team has adequate practice in how to respond to these incidents.

Red teaming

From time to time, it is valuable to test the actual effectiveness of the facility security and its teams' preparedness. Red teaming is the act of performing announced (or unannounced) attacks on the facility to see how teams respond, and eradicate failures and other problems that may occur should an actual attack occur. This greatly prepares teams, but is only effective once security measures are put in place.

Contingency plans

The typical course of business itself can set events in motion that cause security to fail. For instance, the turnover of key personnel, the rapid reduction or reallocation of budgets, lawsuits, audits, etc., can all strain personnel resources which can adversely affect the security posture of the entire infrastructure. It is important to have these contingencies identified, and plans for the events in which they occur.

Conclusion

Our research and findings outlined here cover a wide range of typically overlooked, yet serious security theory, practice, and actionable recommendations. We hope that by disseminating this information to the public, the community can benefit from this research and prepare the healthcare industry for the security threats of today and the future.

Correct mission	**x**	**Modern approach**	**=**	**Success**
Focusing on patient records ***and*** *patient health*		*Addressing advanced threats*		*Patients better protected*

In this paper we've discussed the factors driving healthcare security in the wrong direction, and the means for its course correction. We introduce a new threat model for how healthcare security *should* view security, as well as information to justify this new model. We overview a variety of attack scenarios shown possible through our research, and discuss the ramifications should attacks of this kind be carried out, and the most common design and implementation flaws that lead to these attacks being possible. Lastly, we provide a blueprint for hospitals to follow in order to most effectively migrate from an insecure posture toward the most secure.

We hope that our efforts here can be an aide, if nothing more than a starting point, for security and healthcare professionals and other interested parties to move the security of the industry and their individual organizations in the right direction. As technology and the industry evolves, we welcome and anticipate great changes to these suggestions here, and plan to continue researching new and better ways to reach the end goal: protecting patients from cyber threats.

The industry would benefit greatly from research that addresses the following areas: reshaping hospital budgets so that they can most effectively account for proper security initiatives, addressing security issues found in active medical devices and other primary attack surfaces that directly interface with patients, how to reorganize hospitals to better serve security by granting the appropriate supervision of digital assets to the security personnel, and how to design and implement standards, best practice, or compliance programs that are effective and not counterproductive.

All hospitals that put patient health first should review the blueprint provided here and endeavor to understand and act upon it. Even organizations with a strong security program can always benefit from additional diligence and investigation. We encourage all hospitals to act upon this blueprint, and we welcome feedback for any shortcomings or improvements that may be found. We are also willing to help in any way we can.

Contact Information

For more information regarding this report, please contact Independent Security Evaluators.

Independent Security Evaluators
4901 Springarden Drive
STE 200
Baltimore, MD 21209

phone: 443-270-2296
email: contact@securityevaluators.com
web: www.securityevaluators.com
twitter: @ISEsecurity