

MEDICAL DEVICE THREAT MODELING

This whitepaper describes threat modeling in the context of creating secure medical device systems. International regulators as well as customers are expecting Medical Device Manufacturers to deliver proactively secured devices. This is in part a question of technology, but equally a question of security engineering best practices applied during the product development lifecycle. This includes applying a mature Cybersecurity Risk Assessment methodologies during the Risk Management process. One powerful technique available to engineers is threat modeling, as we will discuss in this whitepaper.

Target Audience:

- Medical Device Manufacturer (MDM) leadership: product development, product engineering, product security, system engineering, quality/design assurance, regulatory affairs.
- Health Delivery Organization (HDO) leadership: information security, clinical systems IT, clinical engineering, HTM

Readers Will:

- Recognize where threat modeling fits into risk management processes of developing medical device systems.
- Improve understanding of threat modeling: when to do it, how it works, and the insights it brings.
- Build a game plan to get their organization started with threat modeling.
- Gain exposure to resources for additional learning.

THREAT MODELING IS EXPECTED IN RISK MANAGEMENT

Threat modeling is an extension to long-standing risk management activities and should be part of cybersecurity risk management when developing medical device software. ISO 14971 is a cornerstone standard to the safety and risk management processes widely used by MDMs. The standard helps MDMs establish procedures to identify and mitigate threats that may result in “physical injury or damage to the health of people, or damage to property or the environment”.

Software is playing a more prominent role both in medical devices (Software in a Medical Device - SiMD) and as a medical device itself (Software as a Medical Device - SaMD). This brings a set of cyber risks beyond the commonly discussed confidentiality, integrity, and availability — most notably, the risk of patient harm. Published in 2016, AAMI TIR57 “Principles for Medical Device Security - Risk Management” attempts to bridge that gap by mapping ISO 14971’s high level process steps (which are focused on managing safety risks) to corresponding steps for managing cybersecurity risks.

RISK MANAGEMENT		
	Safety Related	Cybersecurity Related
Process Guide	IOS 14971	AAMI TIR57 Medical Device & Health IT JSP
Analysis Methods	FMEA/FMECA FTA	STRIDE/PASTA Attack Trees
Scoring Techniques	Probability/Severity Matrix	CVSS OWASP Risk Rating

Table 1: The methods and procedures of safety focused risk management have many parallels when managing cybersecurity risk. Threat modeling plays an important role in modern risk management.

SAFETY VS. CYBERSECURITY ANALOGOUS TERMINOLOGY	
Traditional Safety	Traditional Cyber
Safety: Freedom from unacceptable risk	Safety: Protection from or defense against damage, unauthorized use or modification
Hazard	Threat
Susceptibility	Vulnerability
People, Property, Environment	Asset
Hazard (or Risk) Analysis	(Cyber) Security) Risk Analysis
Misuse (reasonably foreseeable)	Exploit
Sequence of Events	Attack Vector
Hazardous Situation	Event, Incident (potential)
Harm	Incident (occurring), Consequence
Intended Use	Use Case
Probability	Exploitability
Severity	Impact

Safety Risk Management is well established and has a long history. In the late 1950’s, reliability engineers established methods to systematically analyze the failures and effects of faults within complex military systems. Failure Mode Effects Analysis (FMEA) was extended to include criticality in the analysis process (FMECA) and complemented other structured risk analysis methods (such as Fault Tree Analysis - FTA). These prescriptive methods are readily practiced within the medical device risk management domain to identify failure points and the impact of such failure on patient safety.

Just as FMEA, FMECA, and FTA are methods to proactively identify safety risks, threat modeling is a practice to identify cybersecurity risks. Threat modeling frameworks provide organizations with a repeatable way to incorporate key cyber security considerations into their software design and subsequently prevent or mitigate unacceptable compromises to confidentiality, integrity, availability, and safety.

In the same way as MDMs looked across industry to find and embrace tools like FMECA, “threat modeling”, a practice developed in traditional software industry, has been globally advocated within the health industry. Major government agencies have published industry guidance on how to incorporate cybersecurity considerations into the medical device lifecycle (such as United States FDA, Health Canada, Australia TGA, and French ANSM).

These regulatory bodies have established expectations regarding cybersecurity threats and threat modeling. The FDA focuses on considering system level risks and supply chain risks. Health Canada outlines a checklist of general activities a manufacturer should undertake to evaluate and control risk. The TGA asks that MDMs consider cybersecurity practices for manufacturing and the supply chain. ANSM calls for risk analysis, policy for managing and purchasing software components, and verification methods for ensuring there are no vulnerabilities in the software. One difference noted here is that Health Canada does not have language involving the supply chain unlike the other three guidance documents (for more details on regulatory guidance, see MedCrypt’s whitepaper — “[Understanding International Medical Device Cybersecurity Guidance](#)”).

Similarly, publications from private sectors within the health industry have recommended the practice of threat modeling (such as the “[Medical Device and Health IT Joint Security Plan](#)” from Public Health Sector Coordinating Council).

Table 2: Safety risk and cybersecurity risk often use different terms to express comparable concepts. These terms diverge on the fundamentally different assumption that a Safety Hazard is primarily coincidental, while a Cyber Threat is primarily intentional.

KEEP THE PROCESS SIMPLE

Threat modeling is intended to be a systematic and repeatable method of identifying cybersecurity threats that could exploit the weaknesses of your system. As user needs change, intended uses evolve, features are added, and architectures are adjusted, threat modeling needs to be revisited to align with the new or modified aspects of the system. The threat modeling process can be broken down into four key questions. Each question is answered through an associated activity.

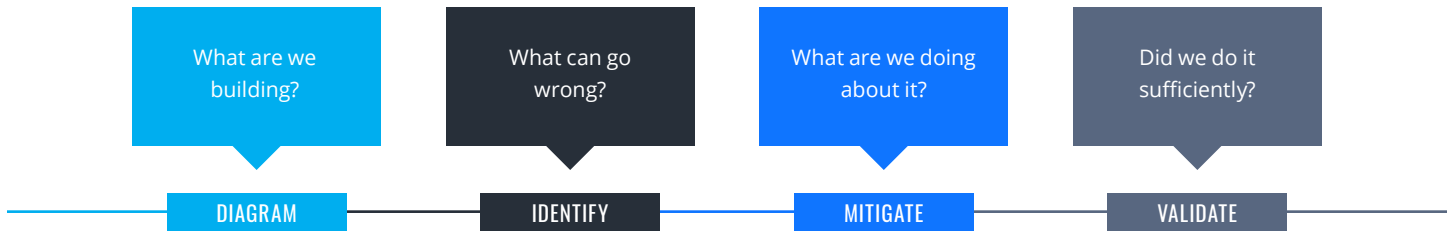


Figure 1: Basic Threat Modeling Steps

MAKE IT A CROSS-FUNCTIONAL ACTIVITY

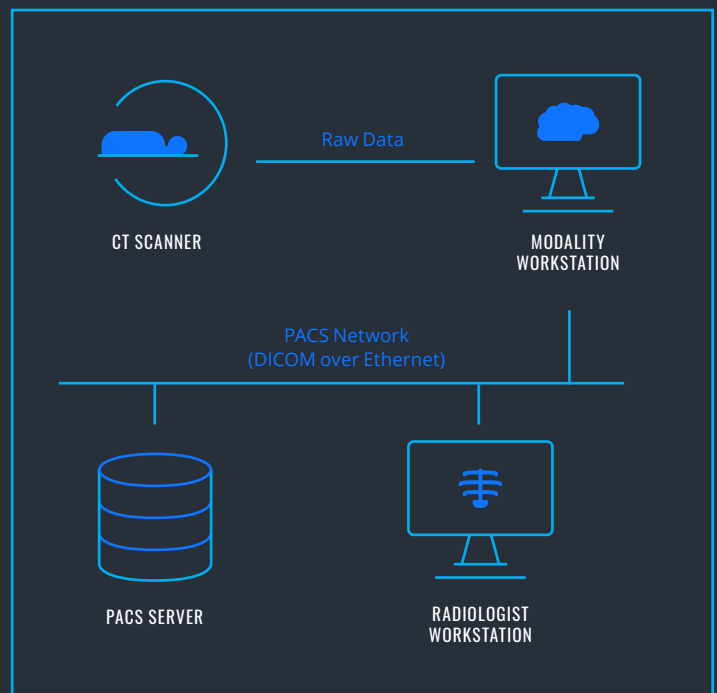
MDMs have a very detailed understanding of the systems they develop and build. The thought of identifying every possible threat in a medical device system can seem overwhelming. Connectivity (network based, serial cables, data carriers, or other) puts devices in a larger ecosystem and creates complex “systems of systems” that have numerous attack surfaces. Attempting to analyze such complexity in a single model is unwieldy and unmaintainable.

In the same way capabilities are logically organized at different levels of detail across multiple requirements documents, threat models should focus on different levels of detail. Start the process from a higher level view-point focused on key use scenarios. As you complete threat modeling of the higher-level system, iterate on modeling sub-systems and components that are most critical to patient safety, data sensitivity, and other important areas. This hierarchical approach to your threat models reduces the maintenance burden as the system evolves. In circumstances of design or use case change, only threat models pertaining to the areas of the system being altered need to be revised. Models of subsystems and components not impacted by the changed do not need to be revisited and updated.

It's likely you already have the system diagrams needed to begin threat modeling activities. HDOs often require MDMs submit security assessments that enable the HDO to understand the risks of deploying your device onto their network. [The Manufacturer Disclosure Statement for Medical Device Security \(MDS2\)](#) is a broadly used assessment template that requires architecture and data-flow diagrams; these diagrams are an excellent starting point for threat modeling.

If you don't have these diagrams already, don't let that stop you from threat modeling. Prepare for deeper threat analysis by whiteboarding different parts of the system and communication paths between those parts. Take a picture of the whiteboard diagram for later conversion into electronic format. The electronic diagrams will not only be helpful in capturing your threat modeling analysis, but they can be reused in HDO device assessments (such as the previously discussed MDS2).

Figure 2: Inform diagram of the medical Imaging system researches hacked to manipulate radiologist's diagnosis of lung cancer. These kind of diagrams provide great starting point for threat modeling.



DIAGRAMMING A REAL-WORLD SCENARIO

In a paper published in 2019, Mirsky et al. ([CT-GAN; Malicious Tampering of 3D Medical Imagery using Deep Learning](#)) document how medical images from a CT scanner can be manipulated in real-time by an attacker to either inject forged evidence for lung cancer, or remove existing evidence. The attack itself is performed by a small hardware device that can easily be disguised to look like an authentic component of the CT system. This device is then able to tamper with medical images generated by the scanner while they are moving from the scanner to the radiologist's workstation. These images are transmitted across the PACS network (Picture Archiving and Communication System) using the DICOM protocol (Digital Imaging and Communications in Medicine).

Figure 3 provides context to the system-of-systems researchers compromised in this scenario. This context diagram is a "user friendly" view and does not employ any formal notation. More formal diagram notations, such as Unified Modeling Language (UML) or Data Flow Diagrams (DFD), make the threat modeling process more digestible. Figure 4 expresses the same systems as Figure 3, but in DFD notation. This makes it easier to understand which parts of the systems are talking to each other, the kind of information they are sending, and the overall flow of data throughout the system.

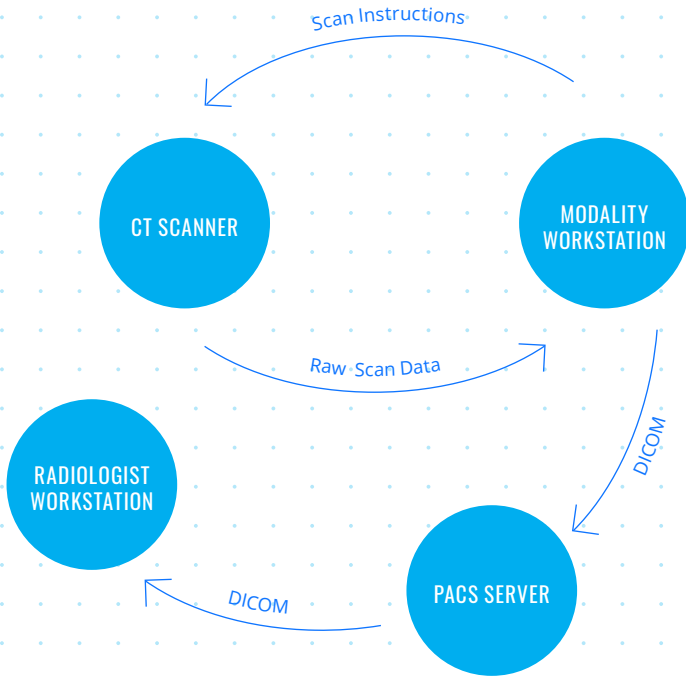


Figure 3: By expressing Figure 3, in DFD notation, we can clearly identify data flows and interactions. These data flows will be analyzed to identify threats and vulnerabilities.

ELICIT THREATS USING STRUCTURED METHODS

The idea of a free-flowing whiteboard hacking session may generate visions of highly energized teams thinking "outside of the box" to produce exotic ways to exploit vulnerabilities. Unfortunately it is difficult to consistently reproduce events like this, and threat modeling needs to be recurring throughout the development process. Existing threat modeling approaches provide structure to guide thinking of threat model contributors while exploring the pertinent constellation of threats.

A good method to elicit threats is flexible enough that it fits your workflow, is repeatable, and facilitates discussion between the right stakeholders. Two examples of such methods are STRIDE and attack trees. STRIDE (as fully described in the book [Threat Modeling: Designing for Security](#)) is a widely applicable method pioneered by Microsoft. It is a mnemonic of different threat categories that can be used to systematically analyze a data flow diagram (e.g., the diagram presented in Figure 4) for threats. Attack trees represent domain specific security expertise. They tell you how you can systematically think about threats against common frameworks, components, or technologies. Figure 5 provides a subset of an attack tree created by security researcher Ivan Ristić which shows multiple ways in which a typical SSL communications setup can be attacked ([SSL attack tree](#)).

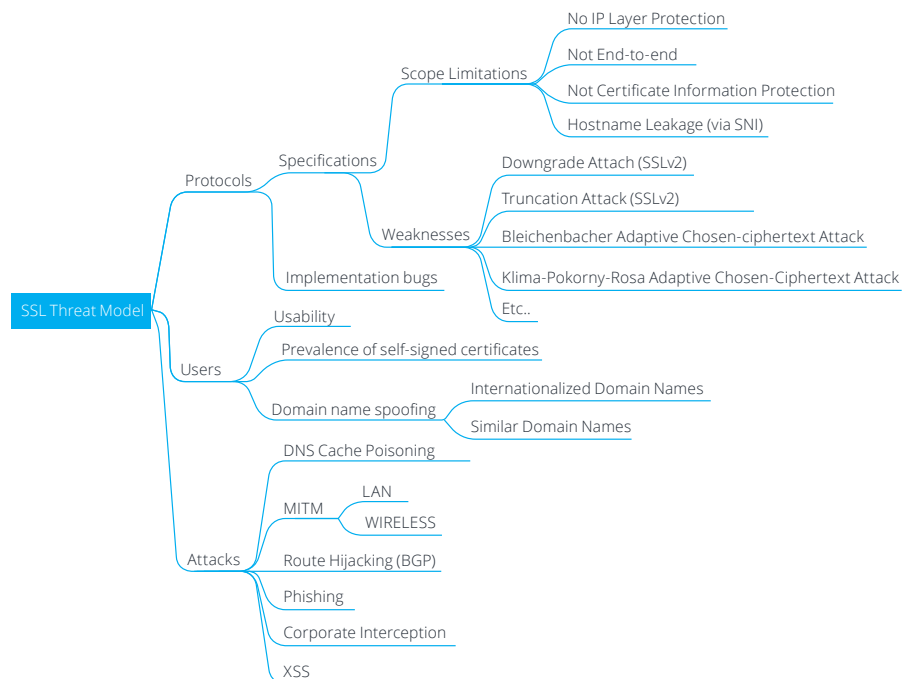


Figure 4: Attack trees provide systemic and visual way to evaluate attack scenarios and understand the attack surface. This example describes way to attack SSL communications; similar trees can be created for medical devices.

APPLYING STRIDE TO THE REAL-WORLD SCENARIO

To illustrate the use of threat elicitation methods, we use the STRIDE approach to analyze communication between the “Modality Workstation” and the “PACS Server” in Figure 4. We are immediately able to make some interesting observations. The analysis tells us that it is important to not only consider protecting the confidentiality of data in transit between the Modality Workstation and the PACS server, but also the integrity. In addition, it is also important to ensure that all parties involved in the data flow from Modality Workstation to PACS server are properly authenticated so that they cannot be impersonated.

Even a basic STRIDE assessment confronts the device manufacturer with many more potential threats that have not been considered in the paper, e.g., is it possible for the radiologist to tamper with scan data to cover up an incorrect diagnosis after the fact? Whether this is technically feasible can then be deferred to an IT security expert, analyzed with threat trees (if applicable threat trees exist for the threat in question), or validated by a scoped security audit.

	Modality Workstation	DICOM Data	PACS Server
Spoofing	...Impersonate the Modality Workstation?	(Not Applicable)	...Impersonate the PACS server?
Tampering	...Tamper with the workstation?	...Manipulate data in transit?	...Tamper with the PACS server?
Repudiation	...Perform actions on the workstation that can later be denied?	(Not applicable)	...Perform actions on the PACS server that can later be denied?
Information disclosure	...Get information out of the workstation without proper authorization	...Intercept data in transit?	...Get information out of the PACS server without proper authorization?
Denial of Service	...Disrupt the workstation?	...Interrupt the data flow?	...Disrupt the PACS server?
Elevation of Privilege	...Perform actions on the workstation that the user should not have access to?	(Not applicable)	...Perform actions on the PACS server that the user should not have access to?

Table 3: The STRIDE methodology Identifies threats as the Modality Workstation sends scan data to the PACS server, The questions in bold map to CT scanner attacks and should be assessed for impact mitigated.

ASSESS RISK LEVELS TO FOCUS MITIGATION EFFORTS

The intention of this exercise is to illustrate there are additional facets of considerations. Every finding identified by threat modeling does not need to be mitigated. However, these findings should serve as the input for a risk assessment, so that the corresponding threats can be ranked. By focusing security efforts on the high priority threats with a large potential business impact, a device manufacturer can both show due diligence and make the overall security process more cost-effective.

Any risk assessment method that is able to quantify technical risks can be used. In theory, a risk assessment method will evaluate each risk in terms of the probability of the event occurring, combined with the impact associated with it occurring. In practice, specific methods exist that decompose the probability and impact factors into specific subcomponents. Probability is decomposed to include ease of exploitation of the weakness, resources and capabilities of anticipated attackers, and so on. Impact is decomposed to include technical impact, business impact, and care delivery impact.

As an illustration, we can analyze the risk apply a well-known scoring methodology, Common Vulnerability Scoring System (CVSS), on the risk that an attacker manipulates scan data in transit to inject false cancer signals. Before the widespread availability of advanced deep fake manipulation techniques, the complexity to pull off an attack whereby DICOM imagery is manipulated on the fly is highly complex, and only has a limited impact on the integrity of that data, as manipulations will not be convincing and would be easy to spot. However, with the rise of deep fake techniques and publicly documented attack strategies, the complexity of the attack goes down while the impact on DICOM image integrity increases. Good assessment methods, such as the CVSS method depicted in Figures 5, easily allow to take this evolving threat context into account.

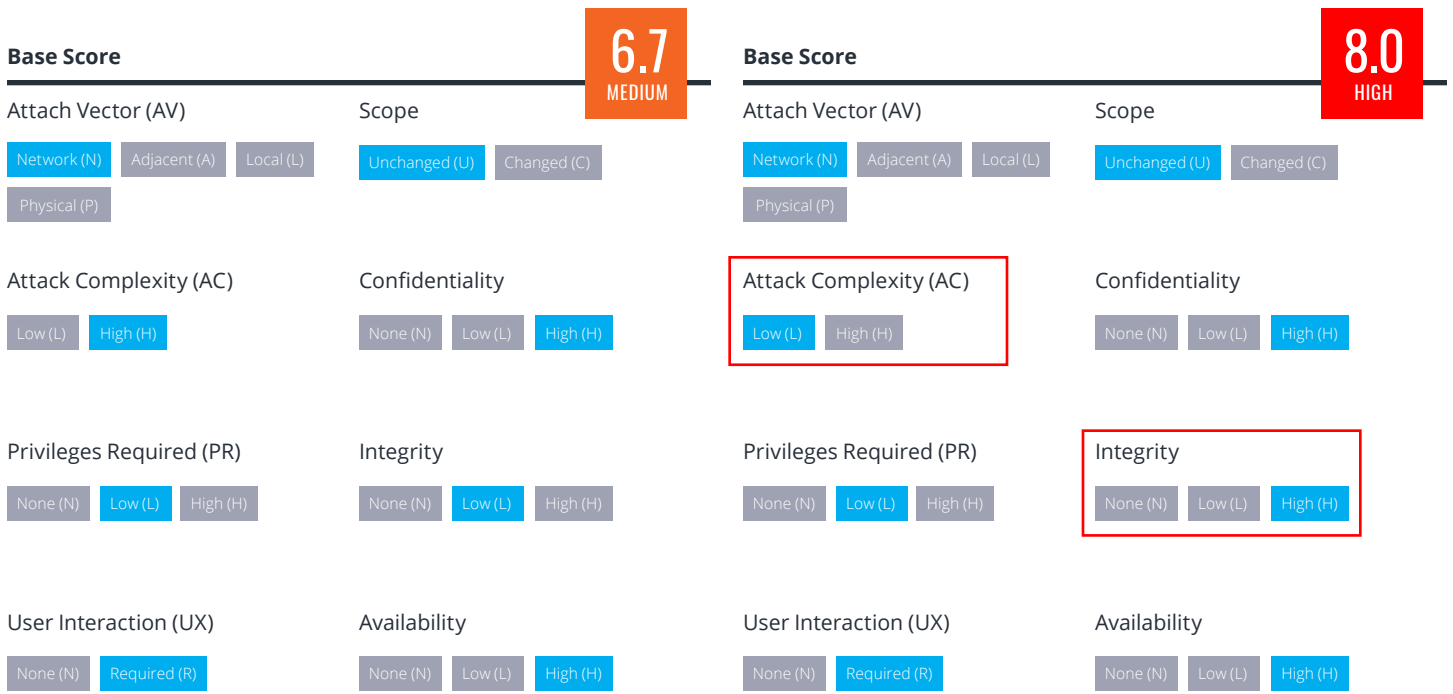


Figure 5: The CVSS scoring calculator uses a formula approach to qualify threat level. Formula inputs should be reviewed on a recurring basis; advancement in technology and changes in device use may alter the values you choose for these formula inputs. The sophistication and prevalence of AI technology increased the score for malicious MRI alteration.

DON'T FORGET TO CONSIDER NON-CLINICAL SCENARIOS

The threat modeling example in this paper focuses on a core clinical use-case for the CT scanner system - transmitting the DICOM images to the PACS server. The Pareto principle infers that 80% of the time a system is being used, only 20% of the capabilities are being exercised. In the case of medical device systems, this rightfully brings attention to the intended uses in a clinical setting. Unfortunately, advanced attackers are clever and will exploit systems in unconventional means; often leveraging backdoors exposed to accommodate capabilities used only 20% of the time. MDMs should also evaluate scenarios outside the clinical setting. Below is a listing of potential threats and example questions an MDM should ask itself:

Manufacturing and assembly processes:

- How is the software/firmware initially loaded on the device by the MDM?
- Is software/firmware pre-loaded on device components provided by upstream suppliers?
- Are secret keys embedded on the device? What is the risk of those secrets being exposed?

Maintenance activities:

- How are software/firmware updates applied to systems being used by patients or HDOs? What precautions are in place to ensure the updates being applied are legitimately from the MDM?
- What are the common maintenance activities? When the device is in "maintenance mode", does it disable safeguards that mitigate threats identified in other scenarios?
- Can a technician remotely access the system for maintenance purposes? If so, how is this done and how are the technician's high-level access privileges protected from compromise?

KEY TAKE-AWAYS AND LOOKING FORWARD

The sooner the better, but never too late

While threat modeling is ideally applied from the early stages of a project, it is never too late to adopt it. When done right, threat modeling is a very flexible technique that scales well in terms of time and effort invested in it; a little goes a long way.

There is more than one way to do it

Multiple approaches to threat modeling exist. In our experience, it is better to start with something basic and improve as you gain experience, rather than postponing until the “perfect” method is defined. The goal of threat modeling is not to simply produce documentation for documentation’s sake; it is to facilitate the right discussions and be explicit about your security posture. These materials can then be shared and evaluated by security experts and regulatory authorities. In this regard, whichever approach helps you to facilitate the right discussions about security has value.

Need continued cross-industry collaboration to tailor the discipline

Although at the time of this writing there has been no documented case of patient harm being caused by a cybersecurity compromise of a medical device, the possibility is widely recognized by MDMs, HDOs, regulators, patients, and even Hollywood movies. Humanity has made many decisions based on unsubstantiated fears so it is reasonable to conclude that if there was a documented event, certain patients would forego treatments on the basis of these concerns. The negative perception of cyber-based risk will not be isolated to a single MDM, it will be felt by the larger industry. Therefore, it is important that MDMs & HDOs recognize the areas of common ground and shared responsibility.

Since the FDA issued its first Premarket Cybersecurity Guidance in 2014, there has been an intentional development of a community around security between HDOs, MDMs, security researchers, and industry organizations. Evidence of this collaborative spirit is seen in an analysis of ICS-CERT cybersecurity disclosures revealing device vendors reported four times as many vulnerabilities per quarter since the FDA released their Postmarket Cybersecurity Guidance in December 2016.

Despite these gains, only a subset of device vendors, representing only a subset of device types, are actively participating in this type of coordinated vulnerability disclosure, indicating that broader adoption of transparency is still lacking in the industry. Although thought leaders have established a path forward, improvement is still required and needs to be across multiple facets of security.

A certain level of information sharing amongst MDMs is necessary to build effective methods that establish a systematic and repeatable framework for securing complex medical device systems. It could be argued that reporting vulnerabilities and sharing details of threat analysis is akin to publishing a playbook for hackers. Thought leaders in cyber security (such as Bruce Schneier) have debunked the myth of “security through obscurity”; the prevalent role of open source software (such as Linux) in mission critical systems provides evidence that being transparent with design and collecting inputs to harden the design, bring results superior to proprietary solutions.

Threat modeling is a generic approach applicable to many domains. Just as FMECA has been tailored to address analysis of specific domains (such as DFMECA and PFMECA), threat modeling techniques can also be tailored. For instance, while STRIDE refers to generic threat categories, certain MDMs have adapted it to include the medical threat categories of “abuse” (asking the question, “can overuse of device therapies cause harm?”) and patient safety (see GE’s presentation at the RSA conference; [Medical Device Threat Modeling with Templates](#)). Similarly, as a generic vulnerability scoring tool, CVSS has been extended to the medical device sector (ref. [MITRE: Rubric for applying CVSS to Medical Devices](#)).

By sharing techniques, attack trees, and models the industry can build consensus on best fit methods and practices to ensure medical devices are truly secure by design and continue to bring innovative ways that improve patient outcomes.

AUTHORS

James Leone, CISSP, CCSK

James enables organizations to rapidly develop and operate secure, cost-effective, and reliable cloud and cloud-to-device systems. After being part of digital initiatives at Sony that proved game-changers for the entertainment industry, he saw the potential digital could bring to the life sciences/healthcare industry. Over a decade later, he continues to passionately apply modern era strategies and tactics to build medical software that is responsible and innovative.

Thomas Heyman, Toreon

Senior security consultant at Toreon, with a PhD in secure software engineering. He has significant experience both teaching and applying threat modeling.

Axel Wirth, CPHIMS, CISSP, HCISPP, AAMIF, FHIMSS, MedCrypt

As Chief Security Strategist, Axel provides strategic vision and industry leadership to MedCrypt and its customers. In this role he helps guide the company in critical security strategy decisions and supports the adoption of leading security technology to the healthcare industry. As an advocate for compliance, privacy, and security — and ultimately patient safety - in healthcare, he draws from over 30 years of international experience in the industry. Wharton School.

Vidya Murthy, MedCrypt

Vidya is fascinated by the impact of cybersecurity on the healthcare space. Beginning her career in consulting, she realized a passion for healthcare and worked for global medical device manufacturer Becton Dickinson. She has since joined MedCrypt, a company focused on bringing cybersecurity leading practices to medical device manufacturers. Vidya holds an MBA from the Wharton School.

ADDITIONAL REFERENCES

CVSS Scoring Calculator:

<https://www.first.org/cvss/calculator/3.0>

DFD diagrams generated using Microsoft's Threat Modeling Tool (<https://docs.microsoft.com/en-us/azure/security/develop/threat-modeling-tool>) and the Medical Device Templates.

<https://github.com/microsoft/threat-modeling-templates>

Threat Modeling: Designing for Security, February 17, 2014, Adam Shostack, available at:

<https://www.amazon.com/Threat-Modeling-Designing-Adam-Shostack/dp/1118809998>

Thank you.

Published March, 2020