ICIT | Institute for Critical Infrastructure Technology

Assessing the FDA's Cybersecurity Guidelines for Medical Device Manufacturers

Why Subtle 'Suggestions' May Not Be Enough

Authors:

James Scott, Sr Fellow at the Institute for Critical Infrastructure Technology

Drew Spaniel, Visiting Scholar, Carnegie Mellon University

In practically all matters of cybersecurity within the health sector, the FDA seems to be in a constant state of offering subtle suggestions where regulatory enforcement is needed. The argument against enforcing cybersecurity standards typically centers on the idea that a regulatory presence stifles innovation. Due to the industry's continuous lack of cybersecurity hygiene, malicious EHR exfiltration and exploiting vulnerabilities in healthcare's IoT attack surface continue to be a profitable priority target for hackers.

On January 15, 2016, the United States Food and Drug Administration (FDA) issued the "Draft Guidance for Industry and Food and Drug Administration Staff," advising medical device manufacturers to address cybersecurity "throughout a product's lifecycle, including during the design, development, production, distribution, deployment, and maintenance of the device." The guidelines offer a voluntary framework that organizations can build upon to ensure that their cybersecurity policies, procedures, and strategies proactively address cybersecurity risks in medical devices before the organization, patients, or the public at large, realize financial or reputational harm from the exploitation of an unaddressed vulnerability by an unknown threat actor.

The recommendations build upon NIST's 2014 "Framework for Improving Critical Infrastructure Cybersecurity," which in turn was published in response to President Obama's Executive Order 13636 that advocates the development of a standardized cybersecurity framework that identifies, detects, protects against, responds, and recovers from cybersecurity risk. The recommendations are not regulations. Regulatory frameworks are difficult to develop and enforce because different organizations operate under different constraints. More often than not, the regulations developed are bare minimums, inadequate to the actual threat, because the regulatory body can only enforce according to the maximum capability of the weakest organization. If situational or organizational conditions dictate that an organization needs to adapt their cybersecurity strategy to different criteria or regulations, then the organization can choose not to follow the guidelines issued by the FDA; however, this freedom should not result in the failure to secure medical devices from cyber threats due to knowledgeable disregard, inefficient budget allocation, or lack of trained cybersecurity personnel. The FDA asks that organizations who wish to implement different cybersecurity measures consult with the FDA to verify that the alternative security is sufficient to the value of the data protected and to ensure that the organization has the opportunity to participate in threat information sharing initiatives. Patients who rely on medical devices should not suffer due to the failure of device manufacturers with lackadaisical cybersecurity standards.

The FDA guidelines detail recommendations for identifying, monitoring, and addressing cybersecurity vulnerabilities in medical devices, throughout all stages of the device lifecycle. Cyber threats evolve as malicious adversaries develop new malicious code, attack along novel threat vectors, and target different data and victims. The healthcare sector is at elevated risk to targeted attacks because lack of regulatory device security and the expansive victim pool makes hospitals and healthcare providers tantalizing targets. Healthcare networks tend to be less secure than comparable networks in other critical infrastructure sectors because cybersecurity only recently became a priority. Further, patient data is more valuable than other target data because its invariant nature means that victims can be exploited for a significant amount of time. To address the threat, products should be designed with controls to anticipate vulnerabilities and to mitigate known risks. Organizations must continue to consider how the confidentiality, integrity, and availability of medical devices and patient data can be assured throughout the entire life of the device.

The FDA’s Draft Guidance stresses that organizations should participate in cybersecurity information sharing through an Information Sharing Analysis Organization (ISAO) and that they should develop a native cybersecurity program. ISAOs are inclusive to any organizations that wish to participate, actionable upon the information received, transparent in how information is communicated among partners, trusted in that the information shared is verified, and the shared information is shielded from public release otherwise required by the Freedom of Information Act or State Sunshine Laws. Further, if the information satisfies the requirements of the Critical Infrastructure Act of 2002, then it is exempt from regulatory use and civil litigation pleas. ISAOs aggregate cybersecurity threat information from the public and private sector so that threats are better understood and anticipated. A dedicated cybersecurity team consists of trained experts and is developed according to a structured and systematic risk management framework. The team ensures that the organization has the best cybersecurity posture it can afford and that any cybersecurity incidents are properly managed. The cybersecurity team develops internal policies and communication channels within the organization to ensure that the organization follows or exceeds the cybersecurity best practices offered by the FDA, NIST, or the information security community at large. The team distributes information to the ISAO, requests information in turn, and adjusts the organization’s security posture in response to current risk. The FDA guidelines, a comprehensive risk assessment of the organization, and the information from the ISAO are used to draft a series of incident response plans, which dictate decision making in the event of crisis. The incident response plans help to stymie the impact of successful attacks and to prevent rash decision making in times of stress.

The FDA cybersecurity guidelines for post-market medical devices are legally non-binding. Corporations, government entities, or individuals cannot call for legal action in response for failure to adhere to the recommendations proposed by the FDA in January 2016. Organizations will still be held accountable to comply with cybersecurity measures set by regulations such as HIPAA or FISMA. The main liability that healthcare organizations face as a result of failure to secure their data is harm to their reputation or the reputation of their partner organizations. In critical infrastructure sectors, the two weakest links are unaware personnel and insecure third party networks. Most major breaches, such as Target and OPM, are the result of a combination of the two attack vectors. In the healthcare sector, it is especially important for medical device manufacturers to design with security in mind and to “monitor, identify, and address cybersecurity vulnerabilities and exploits as part of their post-market management of medical devices” because medical devices are more frequently networked together to facilitate patient care, than in the past. Consequently, one compromised device can allow attackers to compromise the entire network of devices upon which people’s lives depend. No device manufacturer wants to be the third party responsible for the next Anthem-esque breach (~80 million United States citizens). Even if the media never knows that a major breach resulted from lax cybersecurity maintenance on a single networked device, the healthcare providers breached are sure to discover the source and retire those devices. If the breach is bad enough, the provider, and affiliated organizations, could cease business with the device manufacturer. Depending on the purchase agreement or service level agreement (SLA), the provider or breach victims may seek reparations in court.

The information security community observes that many organizations fail to report known vulnerabilities, exploits, or breaches, because the management fears that the organization will be perceived as incompetent or weak. These decision makers fail to realize that the digital age has brought about a desire for transparency and an active information sharing community. In direct contrast to the former antiquated viewpoint, public and partner organizations appreciate when a company updates or patches its product to prevent exploitation by a malicious actor. Regardless, the FDA guidelines clarify when information must be reported as part of post-market "cybersecurity device hygiene". Only cybersecurity vulnerabilities and exploits that "compromise the essential clinical performance" of the device and have a high likelihood of resulting in serious harm or death as a result of exploitation, must be reported to the FDA. Actions taken to mitigate non-critical vulnerabilities or exploits may be considered "cybersecurity routine updates and patches" under 21 CFR part 806, and do not have to be reported to the FDA. Changes made to improve the quality or performance of a device are considered device enhancements, provided that the changes do not impact the essential critical performance of the device. Generally, routine cybersecurity updates and patches are considered a type of device enhancement. In any case, reporting either category of vulnerability or exploit to the healthcare information sharing community would increase the proactive security and awareness of the community at large.

The FDA guidelines recognize that cybersecurity is a responsibility shared by multiple stakeholders. In the healthcare sector, these stakeholders include medical device manufacturers, healthcare payers, healthcare providers, users, Information Technology (IT) system integrators, Health IT developers, and developers and vendors of IT products not regulated by the FDA. These guidelines, aimed at device manufacturers, are only the first step in clarifying the cybersecurity roles, responsibilities, and best practices for each stakeholder.

The theory of broken windows as applied to cybersecurity dictates that attackers target organizations in the most vulnerable sectors because the profit to risk ratio is significant. As the community at large becomes more resilient, less of the attacks succeed. Attackers are subject to a chilling effect as they realize less profit for their expended resources. Eventually the attackers are either caught or they target a more vulnerable sector. Redirecting attackers to a more susceptible sector is not an ideal solution, but it is a realistic scenario. It is difficult to arrest and prosecute cyber threat actors. Given that healthcare data is significantly more valuable than other types of data, given that some other sectors have more resources to dedicate to cybersecurity initiatives, and given that as many lives are not on line in other sectors, the overall impact on American society will be less if Healthcare is a less targeted sector. As an incentive to medical device manufacturers to adhere to the FDA framework and to participate in information sharing initiatives for the improvement of the healthcare community, the FDA will not enforce certain reporting requirements of the Federal Food, Drug, and Cosmetic Act (FD&C Act).

The FDA Guidelines apply to medical devices that contain software, firmware, or programmable logic, and to software considered a medical device on its own. The recommendations do not apply to experimental or investigational medical devices. The guidelines are framed around the *NIST Framework for Improving Critical Infrastructure Cybersecurity* principles of identify, Protect/Detect, and Protect/Respond/Recover. It is the responsibility of the information security team to iterate through these stages and adapt the organizational strategy according to the threat landscape and available resources. It is the responsibility of the executive board to listen to the recommendations of their information security team.

The Identify stage begins when the manufacturer defines the essential clinical performance of the device. Essential clinical performance is the level of performance and function necessary to achieve freedom from unacceptable clinical risk. Defining essential clinical performance depends upon the requirements necessary to achieve device safety and effectiveness. Compromise of the essential critical performance results in the ability to do harm. As part of a comprehensive risk assessment, the information security team and executive board of device manufacturers should define the essential critical performance of a device, they should map the severity of different consequences resulting from the exploitation of a vulnerability, and they should decide the acceptable risk criteria for the device. This information allows executive decision makers the ability to holistically evaluate investments in cybersecurity in comparison to possible outcomes. Acceptable and cost effective mitigation strategies vary greatly on device purpose and capability. It may not be as important to secure a digital thermometer, as it is to secure an MRI machine. The identification stage ensures that limited cybersecurity budgets are efficiently allocated to address the threat landscape and insulate the organization and its customers from harm.

The identification stage concludes with the identification of cybersecurity signals and any necessary action upon those signals. Cybersecurity signals are any indicators of vulnerability or compromise reported in complaints, returned products, service records, internal investigations, post-market surveillance, ISAOs, CERTs, security researchers, or other critical infrastructure sectors. Device manufacturers are already required to analyze signals according to 21 CFR 820.100; however, many manufacturers lack a consistent and comprehensive process. Manufacturers should establish a dependable and repeatable strategy to detect and investigate cybersecurity signals in their devices. This could include participating in an ISAO, incorporating detection mechanisms into devices, or investigating all signals, regardless of source, as vulnerabilities until cyber-forensics proves otherwise.

Detected vulnerabilities should be characterized and assessed to inform triage remediation activities in the organization and in the healthcare community. One way to characterize vulnerabilities is according to their exploitability according to: remote exploitability, attack complexity, threat privileges, actions required by intended user, exploit code maturity, and report confidence. A scoring system, such as the "Common Vulnerability Scoring System (CVSS)," might assist categorization. The next step of the risk assessment is to model threats to each device that is still on the market. Since risk analysis is iterative, the threat model for each device will change until the device is retired. The goal of risk analysis and threat modeling at this stage is to develop strategies to triage vulnerabilities in the least amount of time. Threat modelling typically consists of identifying attack objectives, threat vectors, and vulnerabilities and then identifying countermeasures to prevent or mitigate the threats to the system. As a result, it is useful for the personnel of the information security team to be able to think like a threat actor. The result will be a matrix of attack vectors to devices. In many cases, it will be useful for device manufacturers to analyze threat sources and to postulate their identity, intent, targeting method, tools, techniques, and procedures. It may help to research active threats who target the healthcare sector and other critical infrastructure sectors. Such an analysis identifies new or unconsidered adversaries and it helps to model unexplored attacks. The likelihood, severity, and impacts of each scenario should be evaluated according to a predefined set of criteria and then the results should be documented in a concise summary report per scenario. Unlike other target systems, many medical devices are incapable of detecting threat activity. They may be reliant on additional devices or network monitoring as a result. In the future, it may be useful to incorporate threat detection mechanisms into the design of novel devices. For current devices, the information security team should draft a list of indicators of compromise and the relevant detection mechanisms. This can help network security engineers set system rules and flags to detect malicious activity before harm is realized. For instance, it may be suspicious if a machine that dispenses pain medication were receiving TCP packets from an outside connection. If the healthcare provider IT staff knows to build a firewall rule blocking the connection or an IPS/IDS rule to bring attention to the connection, then a patient might be saved from a targeted attack that could have forced legal ramifications on the Healthcare provider (should the incident be recognized or reported). Signals identified in the risk identification stage or the risk analysis process should be analyzed horizontally (across all devices in the manufacturer's portfolio) and vertically (according to specific components in the device or network). This applies the already conducted risk assessment to similar devices in the product line, to products in development, and to products entering the market.

The third stage revolves around protecting devices, responding to signals/threats, and recovering from the exploitation of a vulnerability. The FDA recommends incorporating device-based features into the design phase as the primary mechanism to mitigate risk to the essential critical performance of the device. Additionally, (and for devices that lack on-board controls) device manufactures should implement compensating controls. A compensating control is an external safeguard or countermeasure, which provides supplementary or comparable cyber protection to the device and its user. One example of a compensating measure might be removing a vulnerable device from the network to ensure that it fulfills its essential critical function. Another example might be a secondary physical control, such as a restricting valve on a pain medication line, which limits the impact of a successful attack. Overall, compensating controls ensure a defense-in-depth approach to device security. They also enable device manufacturers to manage older devices that are too costly or too old to secure according to modern requirements. On the other hand, compensating controls also increase the responsibility and accountability of device manufacturers. It is the responsibility of device manufacturers to notify users of official fixes, temporary fixes, and work-arounds. The FDA recommends manufacturers release compensating control information and disclose vulnerability information according to ISO/IEC 29147:2014 *Information Technology - Security Techniques – Vulnerability Disclosure*.

The final stage of the FDA's suggested framework focuses on mitigating the risk to Essential Clinical Performance. The information security team should aggregate the available information to assess whether the risk to essential critical performance possible by the exploitation of a vulnerability has been adequately addressed by device features or compensating controls. If the residual risk levels remain unacceptable according to the adopted risk management framework and the risk appetite of the organization, then further action must be taken. Ideally, future devices should be designed with internal mechanisms to ensure and validate whether updates, mitigation strategies, and remediation efforts are effective and comprehensive. Otherwise, a compensating control must be included in the process. In any case, mitigation strategies should be proportional to the risk presented and the value of the asset.

Remediation plans should account for a before/ after residual risk evaluation, a risk/ benefit analysis of the solution (because controls can introduce new risk vectors), and steps to mitigate the cascading impacts that result from the remediation measures. When a breach is discovered, the victim organization should notify the relevant investigative agencies before launching mitigation and remediation steps, in order to preserve any cyber-forensic indicators. Afterward, the adversary should be removed from the network and then the breach should be disclosed to victims and the public. Device manufacturers are required to notify users and victims within 30 days of the discovery of a vulnerability in a medical device; however, investigations take time. Consequently, a distinct incident response plan, managed by a qualified information security team, is paramount to remaining compliant to the FDA guidelines within that timeframe. Depending upon the device manufacturer and the affected product, disclosure outside of the initial 30-days may violate regulations. Further, if users or patients are harmed by the compromised device before the device manufacturer discovers and responds to the incident, then the organization, partner organizations, or customers may suffer additional harm.

The medical device community is compliance-oriented. Currently, healthcare device manufacturers and healthcare providers have the ability to ignore the FDA's recommendations. However, it is in the best interest of each organization and the community at large if the target audience pays attention to the FDA's underlying message to adopt a comprehensive risk-based cybersecurity program. Interested stakeholders have 90 days from the January release of the guidelines to submit comments and suggestions to the FDA about the guidelines. It may be beneficial to healthcare providers, healthcare payers, and legislators to petition the FDA to make the guidelines regulatory. Otherwise, medical device manufacturers could ignore the guidelines altogether. Sadly, a survey conducted by Veracode and HIMSS as part of Veracode's "State of Web and Mobile Applications Security in Healthcare" revealed that only 14 of the 200 healthcare IT organizations surveyed believed that IoT devices – medical devices, POS devices, etc.—were a top security threat. In all fairness, these organizations prioritized combating exploitation of applications (28 percent) and preventing social engineering attacks on personnel and insider threats (26 percent). Besides the security incentive, savvy device manufacturers can adopt the guidelines to gain long-term competitive advantage over their rivals. The medical device market is flush with similar products from numerous manufacturers. No rational buyer would purchase an untrusted device when a comparable product comes with assurance of greater device integrity. Compliance with the FDA guidelines provides a demonstrative differentiating factor that compliant device manufacturers can market to healthcare providers and end users.

The cyber threat is real and bad actors are continuously evolving in both stealth and sophistication. Regardless of how medical device manufacturers and healthcare providers receive the guidelines, the FDA has clearly indicated that medical device cyber security is a priority. The healthcare community should note the gesture and take the initiative to assess their own networks and improve their cybersecurity. The healthcare community has until April 21, 2016 to submit comments on the guidelines to the FDA.

Links

Website: www.icitech.org

https://twitter.com/ICITorg

https://www.linkedin.com/company/institute-for-critical-infrastructure-technology-icit-

https://www.facebook.com/ICITorg

Sources:

Date Protection Report:

http://www.dataprotectionreport.com/2016/02/fda-issues-guidance-on-medical-device-cybersecurity-and-interoperable-medical-devices/

United States Food and Drug Administration (FDA):

http://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm481968.htm

Tech News World:

http://www.technewsworld.com/story/83042.html

Threat Post:

https://threatpost.com/fda-issues-guidelines-on-medical-device-cybersecurity/115915/