Title: Cybersecurity and Medical Devices: A Practical Guide for Cardiac Electrophysiologists

Short Title: Cybersecurity and Devices

Benjamin Ransford, PhD[1], Daniel B. Kramer, MD, MPH[2], Denis Foo Kune, PhD[1]; Julio Auto de Medeiros[3]; Chen Yan[4]; Wenyuan Xu, PhD[5]; Thomas Crawford, MD[6]; and Kevin Fu, PhD[1]

[1] Virta Laboratories, Inc., Ann Arbor, MI
[2] Richard A. and Susan F. Smith Center for Outcomes Research in Cardiology, Division of Cardiology, Beth Israel Deaconess Medical Center, Boston, MA
[3] Office of Information Security, Mayo Clinic, Rochester, MN
[4] Zhejiang University, Hangzhou, China
[5] Department of Computer Science & Engineering, University of South Carolina, Columbia, SC
[6] Department of Internal Medicine, Frankel Cardiovascular Center, University of Michigan Health System, Ann Arbor, MI

Address for Correspondence:
Kevin Fu, PhD, Virta Laboratories, Inc., kevin@virtalabs.com; 1327 Jones Drive, Suite 110, Ann Arbor MI 48105. (p) 734-764-1688

Funding: Dr. Kramer is supported by the Greenwall Faculty Scholars Program, and is a consultant to Circulatory Systems Advisory Panel of the Food and Drug Administration. Virta Laboratories, Inc. is supported by SBIR Phase II Grant No. 1555816 from the National Science Foundation. There are no other financial or commercial financial conflicts of interest related to the study topic to report.

Condensed Abstract

Medical devices increasingly depend on software. While this expands the ability of devices to perform key therapeutic and diagnostic functions, reliance on software inevitably causes exposure to hazards of security vulnerabilities. This article uses a recent high-profile case example to outline a proactive approach to security awareness that incorporates a scientific, risk-based analysis of security concerns that supports ongoing discussions with patients about their medical devices.

Keywords: Pacemakers; Implantable Cardioverter-Defibrillators

Background

Although widespread computerization of medical care enables new innovations and improves patient outcomes, the healthcare industry struggles with cybersecurity. With volumes of patient data and increasing dependence on software for lifesaving therapies, providers worry that security gaps could interrupt care, allow for identity theft, or harm patients. These concerns are particularly stark for the medical device industry. Responses to cybersecurity challenges in the last decade have been inconsistent, with some progressive manufacturers developing in-house security programs before problems occur, while others are less proactive about potential security vulnerabilities in their products.

Clinicians and patients remain relatively uninformed about the methods for evaluating security risks, and thus vulnerable to misinformation. Interpreting the results of security research can be challenging for physicians and providers, leaving several questions unanswered: Should security vulnerability reports influence prescribing practices or otherwise affect patient care? What evidentiary standards are appropriate? How do vulnerabilities relate to attacks and patient safety, and how should the likelihood of real compromise be estimated? What is an appropriate response to patients who ask about security vulnerabilities they have seen in the news? This article uses a recently reported case of a potential security vulnerability to: (1) provide an overview of cybersecurity research methods as applied to medical devices; and (2) demonstrate these methods as applied to a specific suspected security threat.

St. Jude Medical and Muddy Waters: Background

Unlike the medical device industry, no single regulatory body oversees software cybersecurity as a whole, and problem reports often originate with independent researchers. The accepted best practice among security researchers is “coordinated disclosure,” wherein a researcher

notifies a software maker and confirms a remediation in advance of public announcements. These reports are similar to “safety communications” issued by the Food and Drug Administration (FDA), but without specific regulatory oversight.

Occasionally researchers sidestep coordinated disclosure. In August 2016, a hybrid market research–vulnerability report written by a hedge fund in concert with a team of security researchers alleged vulnerabilities in St. Jude Medical Merlin™ compatible cardiac implantable electronic devices (CIEDs) and the ecosystem of devices supporting those CIEDs.(1) Regulators, the manufacturer, providers, investors, and other security researchers scrambled to respond.

According to the report, security researchers at a company called MedSec studied St. Jude Medical’s Merlin product line and found several ways in which they believed the products were vulnerable to malicious intrusions. Instead of first contacting the manufacturer, they chose to provide their findings to a hedge fund, Muddy Waters LLC, which publicly announced that it held a short position on St. Jude Medical’s stock — a wager that the stock would decline in value. The report stated that the security researchers would share in profits from the short sale.

The hedge fund’s report alleged two types of attacks: (1) a “crash” attack purportedly causing the CIED to “stop working,” and (2) a “battery drain” attack that could reduce the CIED’s time until replacement. The report asserted that the attacks could be “executed on a very large scale” and “highly likely could be exploited for numerous other types of attacks,” further claiming that “the product safety issues [...] offer unnecessary health risks and should receive serious notice among hospitals, physicians and cardiac patients.(1) A physician’s open letter on University of Chicago stationery at the end of the report stated that he had stopped implanting the affected devices and had recommended patients disconnect their Merlin@home units despite consensus guidance from the Heart Rhythm Society on benefits to patients from remote monitoring (2);

elsewhere, the report acknowledged that the doctor was a board member of the security researchers' firm.(1)

In reference to the "crash" attack, the report described a loss of radio connectivity with the CIED after sending it undisclosed radio traffic. A companion video showed a failed attempt of an operator to program the CIED after sending it the undisclosed radio traffic for several hours.(3) The report also referred to "rapid pacing" correlated with a "crash" attack and presented a screenshot showing a Merlin programmer display as evidence of malfunction.

St. Jude Medical responded by disputing the specific vulnerabilities and the impact of vulnerabilities that might be found. The FDA issued a safety communication in January 2017 that outlined the clinical concerns.(4) Importantly, this communication emphasized that there were no reported patient harms identified, and characterized the potential vulnerability and the software patch developed and validated by St. Jude Medical.(5) Notably, the communication as well as the manufacturer's own guidance recommend keeping patients' remote monitoring systems active to allow for software updates and patches – guidance that contrasts starkly with the physician letter included in the initial security report.

The FDA's involvement deepened after its January 2017 safety communication. An April 2017 warning letter from the FDA to Abbott, which had recently completed an acquisition of St. Jude Medical, stated that St. Jude Medical had "failed to accurately incorporate the findings of a third-party assessment" of cybersecurity risk from 2014 – which the Muddy Waters report also stated – and that St. Jude Medical had failed to follow its own Corrective and Preventive Action (CAPA) process when responding to the Muddy Waters report.(6)

Standards of Evidence in Security Research

How should this security report be viewed by the clinical community? The currency of security research outside healthcare is the *proof of concept*, usually executable program code embodying an *exploit* that takes advantage of a vulnerability in a reproducible way. Unlike medical research, nearly all security research concerns human-made systems that perform *deterministically* and *identically* across every running instance, i.e. a proof of concept will either work or it will not. Security researchers often have access to the source code of the systems they study, leading to high-confidence determinations and claims. For this reason, security researchers do not typically conduct randomized trials or even re-run experiments once a proof of concept is developed.

In the case of the August/September Muddy Waters report and videos, in our opinion, the descriptions and demonstrations not only omitted a proof of concept, but they also left room for crucial questions, chiefly: did the purported vulnerabilities affect the *essential clinical performance* of the CIEDs in question? (The report makes a case for omitting a working proof of concept, since it claims that doing so would risk patient harm, but it also makes speculative claims about the impact of the vulnerabilities, as mentioned above.) *Essential clinical performance* refers to the main criterion that the FDA uses to determine the safety impact of a problem report. The shorter question is: were the CIEDs still able to provide their intended therapies during the purported failures?

It is unrealistic to expect generalist security experts to know the intricacies of medical devices. While investigating a potential vulnerability in a medical device, a security researcher should consider collaborating with at least one of three parties: a physician who knows the device well, a regulator, or the manufacturer. If asked to participate in such a study, any of these parties including the physician should expect to help nonmedical researchers understand how the device is typically used and how to test its essential clinical performance.

In the case of the Muddy Waters report, the CIEDs' essential clinical performance would have been straightforward to test with knowledge of where to find test-rig schematics; Figure 1

below shows a standard test rig. In the absence of this structured approach, in our opinion, it is difficult to evaluate the accuracy of the findings of the Muddy Waters report.

"Crash Attack" Outcome: Experimental Model

Using the general framework outlined above, we sought to replicate, as faithfully as possible, the CIEDs' purported failure modes to understand the likely causes in the wider context of essential clinical performance. Our experiments differ from the experiments described in the Muddy Waters report in two important ways. First, we focused on the "crash attack" to the exclusion of the "battery drain attack." Previous publications articulated and tested the risks of adversarial battery drain in a different manufacturer's CIED.(7, 8) Second, instead of replicating the attack scenario, we sought to replicate the report's results with *legitimate* (i.e., non-adversarial) traffic under a null hypothesis that the purported error conditions were *not* due to malfunction. Thus, we attempted to reproduce the report's "crash attack" outcome—an apparent loss of the ability to communicate with the pacemaker—while testing the pacemaker's therapeutic functions.

We used a St. Jude Merlin programmer version 3650, a new St. Jude Medical Fortify Assura VR implantable cardioverter-defibrillator (ICD) and a new St. Jude Medical Assurity SR 1240 pacemaker furnished on request from St. Jude Medical. We connected pace/sense leads to a custom-made measurement rig according to FDA guidance.(9) To generate radio traffic, we conducted routine interrogations with the unmodified programmer. We monitored the CIED's communication bands using a software radio tuned to the 402–405 MHz MICS band.(10)

The experimental setup, also depicted in Figure 1, was as follows:

- Configure the CIED to pace at 60 bpm and the mode to VVI and confirm that it inhibits pacing in response to a standard simulated cardiac signal.(11)
- Simulate cardiac tissue with a 500Ω resistor per FDA guidance.(9)
- Connect a signal generator (BK Precision 4063) to the sensing input and provide a simulated cardiac signal. Signal characteristics: onset with a linear rise of 2ms followed by a linear fade of 13ms; period of 800ms and peak amplitude of 5mV.
- Connect an oscilloscope to the pacing output via a standard 10MΩ probe and set its display to 900mV/div and 500ms/div (time domain); confirm that, in the absence of a simulated cardiac signal; the pacemaker emits 60 bpm pacing pulses.

Assurity SR Pacemaker

A video released along with the Muddy Waters report claimed a "crash" condition in which an Assurity SR pacemaker was no longer available for telemetry or interrogation after a certain duration of undisclosed radio traffic. We successfully replicated these "crash" conditions against an Assurity SR pacemaker — but *without* affecting its essential clinical performance.

First, the report and video claimed that radio telemetry — the mechanism by which the CIED can communicate with a bedside monitor for active monitoring of the patient — became unavailable after an undisclosed amount of radio traffic. We posited that a sufficient amount of clinically uninteresting interrogation radio traffic would trigger a battery-saving mechanism in the CIED. We initiated a series of twelve radio telemetry connections at regular intervals over a two-hour period. After sending this clinically unusual amount of innocuous traffic, we confirmed via wand interrogation (which uses an inductive near-field channel) that the pacemaker had stopped sending radio telemetry.

Second, the video suggested that the pacemaker had stopped working altogether because the researchers were unable to interrogate it using the programmer wand when the pacemaker was directly atop the programmer. We placed the Assurity SR pacemaker in the same position on the programmer as depicted in the video (on the surface of the open programmer, near the handle) and confirmed that the programmer failed to establish communication for wand interrogation over several attempts. Moving the pacemaker to a different location (a wooden table next to the programmer) allowed normal wand interrogation to be reestablished.

However, during and after both of these purported "crash" conditions in our experiments, we confirmed with the test circuit that the pacemaker correctly emitted pacing pulses at the programmed setting of 60 bpm and correctly inhibited pacing in response to the test cardiac signal. Thus, while telemetry could be inhibited, there was ***no apparent impact on the essential pacing function of the device***.

Fortify Assura VR ICD

The report offers a screenshot of a programmer showing alerts as evidence of an "apparent malfunction" of an Assura ICD they had subjected to a "crash attack." To reproduce the same condition without causing a malfunction, we connected the ICD's ventricular port to a set of resistors following standard practice for testing CIED connections(11), and left the ICD undisturbed for roughly three hours. The "red error messages" ("VS2" markers) in the programmer's screenshot indicated that the ICD sensed ventricular beats, a normal response to electrical noise according to the Merlin PCS help manual. The programmer raised three alerts, as in the screenshot, related to the disconnected lead. The screen displayed an average ventricular rhythm of 162 bpm when the lead was disconnected, suggesting that the screenshot evidence provided by Muddy Waters was not specific to any particular abnormality or device malfunction.

Summary of Experimental Model

With knowledge of clinical testing practices, it was relatively straightforward for us to overcome the experimental shortcomings of the Muddy Waters report, but such knowledge is likely out of scope for most non-specialist security researchers. The key lesson for providers and practitioners reading medical-device security research is that it should be approached in the context of essential clinical performance. Speculative claims offered without basic testing of therapeutic functions should be evaluated *after* the establishment of a clinical baseline. More generally, providers and patients should be aware that there are in fact established standards both for rigorously evaluating and reporting security concerns in medical devices — and we found no evidence that the short sellers followed FDA guidance on reporting cybersecurity problems, which are reiterated clearly in the relevant safety communication.(4)

St. Jude Medical sued the hedge fund, the security researchers, and the collaborating physician named in the report. In the intervening months, the hedge fund and researchers released follow-up videos purportedly demonstrating more vulnerabilities, as well as a website and a report by independent experts offering further analysis of the Merlin products.(12) These follow-up materials met a higher evidentiary standard, with clearer demonstration of experimental methods, but they focused on other vulnerabilities and did not directly support the claims made in the original report (e.g., analysis of the "crash attack" was inconclusive).

What Should Physicians Tell Patients?

Patients are right to wonder about the security of computing devices in or on their bodies, especially when they depend on those devices for lifesaving therapies. They are also right to wonder

whether devices on their home networks, such as home telemetry receivers, introduce security or privacy risks.

The correct answer to these questions is that, like any therapeutic product, no software or hardware medical device is entirely without risk, but clinically proven therapeutic benefits should be weighed more heavily than clinically unproven security hazards when deciding whether to recommend a therapeutic or diagnostic device.(4) The short answer for patients is that they are almost certainly better off with their therapeutic devices than without them. More generally, FDA approval indicates that the manufacturer has provided reasonable assurance of safety and effectiveness, meaning the therapeutic benefits and potential harms have been evaluated thoroughly in the context of the device type and intended indications. However, while medical device software is reviewed to ensure that it is developed and validated using the appropriate practices, FDA review is not an exhaustive, line-by-line examination, and security threats often remain theoretical. Medical device engineering involves consideration and testing of the vast majority of relevant failure modes, especially those related to essential clinical performance. The industry trend is toward better cybersecurity, with new publications such as the FDA's postmarket cybersecurity guidance(13) and AAMI's Technical Information Report 57 recommending specific actions manufacturers can take to improve product security. We are not aware of any reported incidents of targeted medical-device hacks causing patient harm. A prior study evaluating public FDA databases of adverse events and recalls noted that while computing and software capabilities were common among affected devices, specific security and privacy risks were not identified.(14) However, the security hazards to medical devices should not be ignored, and legacy systems relying on passive adverse event collection may not be well suited to identifying security risks.(14) Thus, individual security reports must be carefully validated to determine clinical impact. It behooves physicians to stay abreast of software-related safety communications along with all other FDA

communications, but overreacting to medical device security claims unvetted by FDA or recognized experts should be discouraged. Individual vulnerabilities, if important enough to warrant extra scrutiny by the FDA, may trigger further FDA communications in the normal channels.

Conclusion

Security vulnerabilities appear in nearly every software system, including medical devices. When claims of security problems arise, physicians should focus on clinical impacts and demand a coherent standard of evidence. A recent report that, in our opinion, fails to use traditional scientific standards of evidence for security research serves as a cautionary tale for providers, physicians, patients, manufacturers, and security researchers. In particular, security reports on medical devices should take steps to rule out null hypotheses that may represent more plausible alternative causes (e.g., experimental error, electromagnetic interference, ungrounded leads, RF telemetry lockout). As always, alarmism must be tempered by rigor. An elevated temperature could indicate a serious infection, or it could simply indicate that a patient ingested hot coffee. Selecting appropriate null hypotheses for medical device security requires specialized skills and training. The danger of misinterpreting a spurious correlation is that patients may make life decisions that lead to greater risks. Providers can best defend against security incidents, even those that have not yet occurred, by adopting industry standards for cybersecurity and ensuring that procurement practices treat these standards' prescriptions as requirements. Because future medical device security problems could lead to harm and too many reports based on incomplete analysis could foster complacency among providers and manufacturers, we recommend scientific rigor as the best defense to promote cybersecurity as a public good in the best interest of patients. All medical devices need better cybersecurity. However, the key message for providers and patients is that patients prescribed a medical device are far safer with the device than without.

Author Contributions

Ransford contributed to drafting, experiment design, and interpretation of results. Kramer contributed to drafting, revisions, and concepts. Foo Kune conducted experiments and contributed to drafting and interpretation of results. Auto contributed to experimental design and interpretation of results. Yan conducted experiments. Xu contributed guidance to experiments and financial support for Yan. Fu contributed to review of results. Crawford contributed to drafting and review of experimental results.

References

1. Muddy Waters LLC. MW is Short St. Jude Medical (STJ:US). August 25, 2016. Available from: <*http://d.muddywatersresearch.com/wp-content/uploads/2016/08/MW_STJ_08252016_2.pdf*>

2. Slotwiner D, Varma N, Akar JG, Annas G, Beardsall M, Fogel RI, et al. HRS Expert Consensus Statement on remote interrogation and monitoring for cardiovascular implantable electronic devices. Heart Rhythm: the official journal of the Heart Rhythm Society. 2015;12(7):e69-100.

3. Muddy Waters LLC. STJ Pacemaker Crash Attack. Vimeo. [updated August 29, 2016]. Available from: <*https://vimeo.com/180593205*>

4. U.S. Food and Drug Administration. Cybersecurity Vulnerabilities Identified in St. Jude Medical's Implantable Cardiac Devices and Merlin@home Transmitter: FDA Safety Communication. Available from: <*https://www.fda.gov/MedicalDevices/Safety/AlertsandNotices/ucm535843.htm*>. Accessed March 20, 2017.

5. St Jude Medical Press Release, January 9, 2017. St. Jude Medical Announces Cybersecurity Updates. Available from: <*http://media.sjm.com/newsroom/news-releases/news-releases-details/2017/St-Jude-Medical-Announces-Cybersecurity-Updates/default.aspx*> Accessed March 20, 2017.

6. U.S. Food and Drug Administration. Warning Letter to Abbott (St. Jude Medical Inc.) 4/12/2017. Available from: <*https://www.fda.gov/ICECI/EnforcementActions/WarningLetters/2017/ucm552687.htm*> Accessed April 18, 2017.

7. Halperin D, Heydt-Benjamin TS, Fu K, Kohno T, Maisel WH. Security and privacy for implantable medical devices. IEEE pervasive computing. 2008;7(1).

8. Halperin D, Heydt-Benjamin TS, Ransford B, Clark SS, Defend B, Morgan W, et al., editors. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. Security and Privacy, 2008 SP 2008 IEEE Symposium on; 2008: IEEE.

9. Dahms DF. Implantable Pacemaker Testing Guidance. U.S. Food and Drug Administration, editor. 1990.

10. St. Jude Medical. Frequently Asked Questions: Merlin.net Patient Care Network (PCN) 8.0 Q&A. 2015. Available at: <*https://sjm.com/~/media/pro/therapies/merlin-net-patient-care-network-pcn/en/sjmmer091400211pcn80faq.pdf*>

11. Association for the Advancement of Medical Instrumentation. Active implantable medical devices - Electromagnetic compatibility - EMC test protocols for implantable cardiac pacemakers and implantable cardioverter defibrillators. ANSI/AAMI PC69:2007.

12. Livitt CD. Preliminary Expert Report, CASE 0:16-cv-03002-DWF-JSM. MedSec, 2016 25-1.

13. U.S. Food and Drug Administration. Postmarket Management of Cybersecurity in Medical Devices. Guidance for Industry and Food and Drug Administration Staff. Issued December 2016.

Available from: <*https://www.fda.gov/ucm/groups/fdagov-public/@fdagov-meddev-gen/documents/document/ucm482022.pdf*> Accessed March 20, 2017.

14. Kramer DB, Baker M, Ransford B, Molina-Markham A, Stewart Q, Fu K, et al. Security and privacy qualities of medical devices: an analysis of FDA postmarket surveillance. PLoS One. 2012;7(7):e40200.

Figure Legends:

Figure. St Jude Assurity pacemaker pacing and measurement setup.

We simulate cardiac tissue with a 500 Ω resistor across electrodes with a lead in the ventricular IS-1 port, and a 1K Ω resistor from the can to the outer electrode. The oscilloscope (V) measures the voltage between the two electrodes of the lead. One probe from the signal generator is connected to a 4.5K Ω resistor for a 1/10 voltage divider across the electrodes, and the other directly to the ring (anode) of the lead.