Noninvasive Postmarket Security Monitoring for Medical Devices[1]

Benjamin Ransford
Virta Laboratories, Inc.,
Ann Arbor, MI 48105

Denis Foo Kune
Virta Laboratories, Inc.,
Ann Arbor, MI 48105

Ann Gookin
Virta Laboratories, Inc.,
Ann Arbor, MI 48105

Andrew DeOrio
Virta Laboratories, Inc.,
Ann Arbor, MI 48105

1 Background

Software-based medical devices enable fast product-development cycles, constructive information sharing, and configurable therapy delivery, resulting in better patient outcomes overall. An unfortunate drawback is that software is complicated and difficult to maintain correctly. Devices with inadequate software maintenance may pose operational risks to network security and patient safety and privacy. This paper describes a noninvasive approach to medical-device monitoring that can address some of the shortcomings of conventional approaches.

Protecting software-based medical devices from malware infections or network-based mischief is a growing concern for clinical engineers and healthcare information technology (IT) practitioners. Unlike desktop PCs and laptops, medical devices often lack support for antivirus systems or operating-system patches, despite running off-the-shelf operating systems and commercial third-party software. Manufacturers have cited previous regulatory approval as a reason not to support software updates [1], despite the Food and Drug Administration's clarifications to the contrary [2].

Medical devices are often in use for decades in clinical settings, during which time the software they are based on continues to change. For example, Microsoft Windows has undergone four major product revisions since the release of Windows XP in 2001, but new medical devices are shipped with Windows XP as recently as 2012 [3], and anecdotal evidence suggests many more are still in use [4]. Microsoft halted support for Windows XP in early 2014. Even when patches are available, administrators tend to emphasize functionality and efficacy over security and avoid applying patches [4] for fear of breaking systems or voiding warranties.

Without adequate patching, the threat to connected devices increases with time as more vulnerabilities are discovered. Manufacturers have little incentive to retest devices once they are in the field [5], and testing is not guaranteed to catch vulnerabilities. While proactive manufacturers have been steadily improving their design and maintenance processes for new devices to prevent security holes and permit software patches, healthcare IT practitioners are often left with a mess they cannot effectively maintain.

Third-party software also poses challenges for manufacturers and device owners. Devices often ship with commercial or open-source libraries that are maintained separately from a device's main code, often by completely separate teams. Popular libraries that are easy to use become widespread if they add new capabilities to devices. For example, the OpenSSL library for encrypting communications, which is compatible with a range of systems from embedded to server-class, appears in at least 74 different kinds of devices, each with different update mechanisms. A 2014 Internet scan for a particularly high-impact vulnerability called Heartbleed found that 56% of vulnerable devices were embedded systems [6], which are typically more difficult to update than PCs or servers. Modern applications bring together tens of libraries that must all be considered separate sources of potential security problems.

1.1 State-of-the-Art: NIST NCCOE on Infusion Pumps.

The National Institute of Standards and Technology (NIST) recently published a medical-device security use case [5] describing the security risks of a modern medical device, an infusion pump with wireless network connectivity. Network connectivity gives this device a means of adjusting therapy, feeding data into electronic health records (EHRs), and working with centralized management tools. NIST also pointed out the risk of bad actors changing infusion rates to harm patients, a potentially dangerous criminal offense. Independent researchers have confirmed that commercially available devices are vulnerable to such tampering [7]. Malicious hackers may be tempted by a potential foothold into a hospital network [8], a desire to harm specific patients or other mischievous ends [9].

1.2 New Opportunity: Postmarket Monitoring.

The remainder of this brief summarizes an approach to medical-device monitoring that can be added to medical devices in postmarket scenarios, i.e., after they are deployed in clinical settings. Monitoring is a component of security, along with prevention and remediation. In light of the NIST use case, the example in this brief focuses on monitoring infusion pumps for unauthorized dosage changes.

2 Methods

Our system, which draws inspiration from previously published work [10,11], comprises hardware and software dedicated to nonintrusive monitoring. We use machine learning to match a device's activity patterns to previously observed behaviors. The monitoring point for device behaviors is the AC power outlet, a common interface to many medical devices. For many kinds of plugged-in devices, distinct activities on the device correspond to distinct patterns of power consumption. In the case of an infusion pump, the relevant question is whether different infusion rates can be distinguished by examining power consumption.

We studied three commercially available infusion pumps from two manufacturers: two large-volume pumps and a low-volume syringe-type pump. Each runs an off-the-shelf operating system. Our system flags anomalous behavior by comparing behavioral samples against a model derived from a training set. For each pump, we first established a set of training data by recording only the pump's normal behavior with our prototype hardware. We used a proprietary feature-engineering workflow to determine which properties of the collected signals revealed the most about the system's aggregate behavior. With a model derived from the training data, we commenced "live" measurement of the pump to compare behavioral samples to the trained model using the selected features. If a series of new measurements were sufficiently outside the normal range with respect to a sufficient number of features, our detectors automatically deemed the pump's behavior as anomalous.

We qualified "normal" behavior on the two large-volume pumps as an infusion rate of 50 ml/hr, and on the syringe-type pump as an infusion rate of 0.1 ml/hr. Normal depends on the drug being delivered, the patient, and other clinical decision processes, but in keeping with the NIST threat model, we focused on

[1] Accepted and presented at The Design of Medical Devices Conference (DMD2016), April 11–14, 2016 Minneapolis, MN, USA.

DOI: 10.1115/1.4033285

Manuscript received March 1, 2016; final manuscript received March 18, 2016; published online May 12, 2016. Editor: William Durfee.

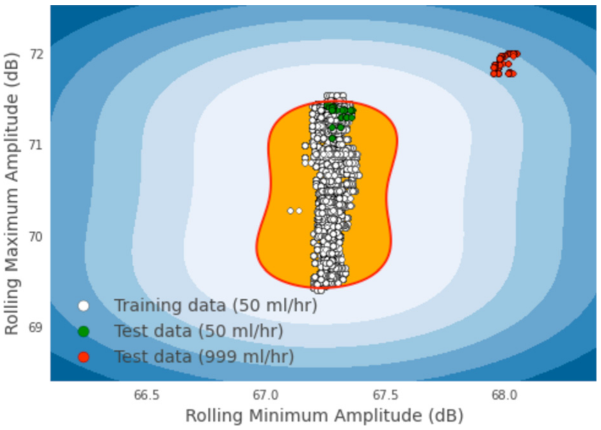

Fig. 1 A two-feature classifier identifies abnormal infusion rates (red points outside boundary) on an off-the-shelf infusion pump

scenarios in which an adversary would be able to change a low infusion rate to an inappropriately high infusion rate. We qualified an "abnormal" rate on the large-volume pumps as 999 ml/hr, the maximum configurable, and on the syringe-type pump as its maximum 20 ml/hr for a simulated morphine sulfate 5 ml/hr syringe.

3 Results

Over all the infusion pumps we tested, our tools automatically generated models that discriminated among the various infusion rates with low rates of false positives (1.0% and lower) and false negatives (close to 0.0%).

Figure 1 is a visualization of an infusion pump anomaly detection model (a one-class support vector machine with two features) that our prototype toolchain produced. The *X*- and *Y*-axes plot the two features used by this model of a high-volume infusion pump. Each dot represents one measurement of the pump's power consumption over 100 s, with respect to two features we used to build the model. Normal data collected from the pump at a 50 ml/hr delivery rate (white dots) were used as training data. Next, we evaluated the model with new measurements of normal 50 ml/hr infusion rates (green dots), which fall inside the red boundary, indicating no anomaly. Finally, we evaluated the model on new measurements that simulated abnormal (attack) activity at a 999 ml/hr rate. These dots fall outside the red boundary, indicating an anomaly. Our error rate on the training data was 1.0%, with 0.0% error for normal test data (false positives) and 0.0% error on abnormal test data (false negatives). The trade-off between false negatives and false positives is configurable in the model. For a second high-volume pump with a similar training and testing regimen, the error on our training data was 0.1% (false positives), and error on the test data was 0.0% (false negatives).

In addition to anomaly detection, we also built a regression model to predict the infusion rate using only features measured from the AC power line; the plot is omitted for space. The input to this model is the *X*-axis feature from Fig. 1, measured over 100 s. The output is the predicted infusion rate, which is a continuous value in contrast to the discrete *normal versus abnormal* output of the anomaly detector. The model suggests that infusion rate is correlated ($R^2 = 0.82$) to the features we measured on the pump's AC power traces.

4 Interpretation

Our preliminary study demonstrates that commercially available infusion pumps' patterns of AC power consumption are correlated with their infusion rates, suggesting that monitoring power at fine granularity is a potentially viable approach to postmarket security maintenance for these medical devices. Postmarket strategies cannot replace coherent patching and update strategies by manufacturers, but they can help decrease healthcare IT practitioners' reliance on manufacturers to provide updates in lockstep with known vulnerabilities in off-the-shelf components.

Acknowledgment

This material is based upon work supported by the National Science Foundation under the SBIR Grant No. 1448440.

References

[1] Baxa, 2012, "Preventing Cyber Attacks," Baxa Corp., Englewood, CO, accessed Oct. 15, 2012, http://blog.secure-medicine.org/2012/06/baxas-non-approved-software-policy.html

[2] U.S. FDA, 2009, "Cybersecurity for Networked Medical Devices is a Shared Responsibility: FDA Safety Reminder," U.S. Food and Drug Administration, Silver Spring, MD, accessed April 8, 2016, http://www.fda.gov/MedicalDevices/Safety/AlertsandNotices/ucm189111.htm

[3] Talbot, D., 2012, "Computer Viruses are 'Rampant' on Medical Devices in Hospitals," MIT Technology Review, Oct. 17 (epub), accessed Nov. 9, 2015.

[4] Kandek, W., 2014, "Windows XP Usage Lower Across Industries," Qualys Inc., Redwood City, CA (epub), accessed Nov. 9, 2015, https://community.qualys.com/blogs/laws-of-vulnerabilities/2014/04/02/windows-xp-usage-lower-across-industries

[5] O'Brien, G., and Khanna, G., 2014, "Wireless Medical Infusion Pumps—Medical Device Security," National Cybersecurity Center of Excellence (NCCoE), U.S. National Institute of Standards and Technology, Gaithersburg, MD, accessed Dec. 18, 2014, http://nccoe.nist.gov/sites/default/files/nccoe/NCCOE_HIT-Medical-Device-Use-Case.pdf

[6] Durumeric, Z., Kasten, J., Adrian, D., Halderman, J. A., Bailey, M., Li, F., Weaver, N., Amann, J., Beekman, J., Payer, M., and Paxson, V., 2014, "The Matter of Heartbleed," Internet Measurement Conference (IMC'14), Vancouver, BC, Canada, Nov. 5–7, pp. 475–488.

[7] Zetter, K., 2015, "Hacker Can Send Fatal Dose to Hospital Drug Pumps," Wired Magazine, Boone, IA, accessed Nov. 9, 2015, http://www.wired.com/2015/06/hackers-can-send-fatal-doses-hospital-drug-pumps/

[8] TrapX Labs, 2015, "Anatomy of an Attack: MedJack (Medical Device Hijack)—Anatomy of an Attack," TrapX Security, Inc., San Mateo, CA, accessed Nov. 9, 2015, http://deceive.trapx.com/rs/929-JEW-675/images/AOA_Report_TrapX_AnatomyOfAttack-MEDJACK.pdf

[9] Williams, P. A., and Woodward, A. J., 2015, "Cybersecurity Vulnerabilities in Medical Devices: A Complex Environment and Multifaceted Problem," Med. Devices, **8**, pp. 305–316.

[10] Clark, S., Mustafa, H., Ransford, B., Sorber, J., Fu, K., and Xu, W., 2013, "Current Events: Identifying Webpages by Tapping the Electrical Outlet," 18th European Symposium on Research in Computer Security (ESORICS), Egham, UK, Sept. 9–13, pp. 700–717.

[11] Clark, S. S., Ransford, B., Rahmati, A., Guineau, S., Sorber, J., Fu, K., and Xu, W., 2013, "WattsUpDoc: Power Side Channels to Nonintrusively Discover Untargeted Malware on Embedded Medical Devices," USENIX Workshop on Health Information Technologies, Washington, DC, Aug. 12.