Published in final edited form as:

IEEE Secur Priv. 2012 ; 10(5): 61–63. doi:10.1109/MSP.2012.128.

Security and Interoperable Medical Device Systems: Part 1

Krishna K. Venkatasubramanian,

Dept. of Computer Science, Worcester Polytechnic Institute, Worcester, MA, USA, kven@wpi.edu

Eugene Y. Vasserman,

Dept. of Computing and Info. Sci., Kansas State University, Manhattan, KS, USA, eyv@ksu.edu

Oleg Sokolsky, and

Dept. of Computer and Info. Science, University of Pennsylvania, Philadelphia, PA, USA,
sokolsky@cis.upenn.edu

Insup Lee

Dept. of Computer and Info. Science, University of Pennsylvania, Philadelphia, PA, USA,
lee@cis.upenn.edu

Medical devices are essential for modern medicine because they can help automate many patient monitoring and management functions. Such devices can be stand-alone or interoperable. Stand-alone devices, by far the most common type, perform monitoring and treatment without directly interacting with other medical devices or equipment.¹ Recently, however, many medical devices have been augmenting their stand-alone operation with considerable communication capabilities, allowing them to interact with other devices. This interoperability offers numerous advantages, including increased safety, usability, and decision support, and a decrease in false alarms and clinicians' cognitive workload.¹

Until now, interoperability has been the domain of large device and systems manufacturers, who require all-or-nothing adoption. That is, all devices must be from the same manufacturer or individually vetted partners. This single-integrator situation is considered safe, owing to these manufacturers' extensive control over interoperating devices, but this solution doesn't scale. Overcoming this problem requires enabling interoperability between different manufacturers' devices but sacrifices control, and has negative economic consequences for traditional device manufacturers. Furthermore, if a failure occurs, the root causes become difficult to trace, which can be problematic for clinical facilities and regulating agencies.

Given the diversity of medical devices that might need to be interconnected, and the structure of economic incentives, the wait for manufacturers to organically evolve interoperability for their devices has already been long. Moreover, regulatory agencies such as the US Food and Drug Administration don't have the mandate to require interoperability. Fortunately, the various stakeholders (manufacturers, clinical facilities, regulating agencies, and so on) are recognizing that the future lies in building genuine interoperability. Consequently, various groups are proactively developing standards that will let devices talk to one another.²⁻⁵

Interoperable medical devices (IMDs) face several threats due to the increased attack surface presented by interoperability and the corresponding infrastructure. Introducing networking and coordination functionalities fundamentally alters medical systems' security properties. Understanding the threats is an important first step in eventually designing security solutions for such systems. Here, in the first part of a two-part article, we provide an overview of the IMD environment and the attacks that can be mounted on it.

The IMD Environment

Because of its flexibility and openness, we use the Medical Device Plug-and- Play Integrated Clinical Environment (ICE) interoperability architecture, as described in the ASTM 2761 standard,² to frame our work. However, the results apply to many other architectures and standards.

Figure 1 demonstrates a simplified view of our ICE-based IMD environment. The *coordinator* is middleware that connects a group of medical devices through a shared network. Legacy devices can interoperate using an *Adapter*. An *alarm system* generates alarms, both medical (for example, related to patient health) and functional (for example, regarding the unavailability of devices, the network, or the coordinator). Individual devices might have additional alarms.

In the rest of this discussion, we assume that if the coordinator fails (for example, in the event of an attack), individual devices independently and automatically enter a non coordinating “offline” safe state and sound their built-in alarm. This assumption is necessary to achieve systems that are safer than current ones. Devices must have a fail-safe mode in case of coordination failure, or patients would face a new risk in the IMD system.

IMD Security

One of the most important issues with such systems of systems is ensuring patient safety, which depends at least partially on the security guarantees offered by the IMDs and connecting infrastructure. If an attacker can force an entity in the IMD environment to deviate from correct behavior, the environment can no longer be considered safe. Furthermore, a compromised device can cause another, otherwise functional, device to perform dangerous tasks. In potentially adversarial situations, such safety concerns are only exacerbated by inter device communication that allows remote access of the entities.

Security is therefore a key requirement for IMDs for two reasons:

They might be deployed in life- critical settings; that is, they might administer treatment, causing changes to the patient’s body, potentially as a result of external directives.

They have access to sensitive health information.

Security attacks on medical devices have thus far been relatively rare, but as IMDs become common, incentives increase to attack them for profit. Moreover, owing partly to laws such as the US Health Insurance Portability and Accountability Act (HIPAA), maintaining security and privacy of patient information is a legal necessity. Recent years have brought increased attention to security vulnerabilities in stand- alone medical devices.⁶ Introduction of interoperability makes devices increasingly connected to and dependent on each other. Because of this increased complexity, the connected devices will likely offer more attack avenues. An adversary needs only to take over the weakest device in the IMD environment to gain a foot- hold. He or she can then reach other devices through the existing trust relationships in the environment.

An Attack Model

Adversaries targeting IMDs come in two basic types. *Passive* attackers can eavesdrop on traffic between IMDs and the coordinator. *Active* attackers can also alter messages, inject traffic, replay old messages, spoof, and ultimately compromise the IMDs’ integrity.

Similarly to Zinaida Benenson and her colleagues,⁷ we designate five classes of attacks on IMD environments: *destroy*, *disturb*, *reprogram*, *denial of service*, and *eavesdrop*. All are active attacks except for eavesdrop. Table 1 illustrates the environment's susceptibility to these attacks.

Destroy

These attacks physically destroy some or all of the components in an interoperability environment, stopping its operation immediately. For example, an attacker could cut an infusion pump tube.

Disturb

These attacks modify the data available to some or all of the entities in the environment to prevent them from operating correctly. Examples include replay and man-in-the-middle attacks.

Reprogram

A special subset of disturb attacks, these attacks modify data or code in a medical device, the coordinator, or the alarm system such that it doesn't perform its designated operation. For example, an attacker could modify an infusion pump's software to deliver extra medication. Reprogramming can be done locally or remotely if a device provides over-the-network programmability.

Denial of Service

These attacks target the network but also affect the devices, coordinator, or alarm system to prevent effective interoperation. For example, an attacker could burn out an infusion pump's motors through overuse, preventing the device from performing the required therapeutic functions.

Eavesdrop

These attacks involve listening in on the IMD environment's network to learn sensitive health information. Because these attacks (unlike the previous ones) don't disrupt system operation, detecting them is difficult.

There's nothing fundamentally new about the attack vectors we presented. However, their use in the context of the coordinating devices and middleware can cause a variety of failures, many of which can't be easily detected because they're silent. In part 2, we'll build on this attack model and demonstrate how adversaries can cause various types of failures in IMD environments, and these failures' security consequences. We'll also introduce the concept of device criticality as a way to assess attacks' potential damage. Finally, we'll conclude with the lessons learned from performing this attack analysis.

Acknowledgments

US National Institutes of Health grant 1U01EB012470-01 and US National Science Foundation award CPS-1035715 partly supported this research.

Biography

Krishna K. Venkatasubramanian is an assistant professor in the Worcester Polytechnic Institute's Department of Computer Science. Contact him at kven@wpi.edu.

Eugene Y. Vasserman is an assistant professor in Kansas State University's Department of Computing and Information Sciences. Contact him at eyv@ksu.edu.

Oleg Sokolsky is a research associate professor in the University of Pennsylvania's Department of Computer and Information Science. Contact him at sokolsky@cis.upenn.edu.

Insup Lee is the Cecilia Fidler Moore Professor of Computer and Information Science at the University of Pennsylvania. Contact him at lee@cis.upenn.edu.

References

1. Arney, D., et al. Proc. 33rd Ann. Int'l Conf. IEEE Eng. in Medicine and Biology Soc.. IEEE; 2011. Biomedical Devices and Systems Security; p. 2376-2379.EMBC 11
2. ASTM F2761 - 09 Medical Devices and Medical Systems—Essential Safety Requirements for Equipment Comprising the Patient-Centric Integrated Clinical Environment (ICE)—Part 1: General Requirements and Conceptual Model. 2009. ASTM F29.21, ASTM Int'l
3. Clarke, M., et al. Proc. 29th Ann. Int'l Conf. IEEE Eng. in Medicine and Biology Soc. IEEE; 2007. Developing a Standard for Personal Health Devices Based on 11073; p. 6174-6176.EMBC 07
4. Introduction to HL7 Standards. Health Level Seven Int'l. 2012. www.hl7.org/implement/standards.
5. Integrating the Healthcare Enterprise. 2012. IHE Int'l www.ihe.net.
6. Halperin, D., et al. Proc. 2008 IEEE Symp. Security and Privacy. IEEE; 2008. Pacemakers and Implantable Cardiac Defibrillators: Software Radio Attacks and Zero- Power Defenses; p. 129-142.
7. Benenson Z, Blaß E, szlig, Freiling FC. Attacker Models for Wire- less Sensor Networks. Information Technology. 2010; 52(6):320–324.

Figure 1.

A simplified generic architecture for interoperable medical devices. The coordinator connects a group of medical devices via a shared network. The alarm system generates alarms for both medical and functional problems.

Table 1

The interoperable-medical-device environment's susceptibility to attacks.

Entity	Attack class				
	Destroy	Reprogram	Disturb	Denial of service	Eavesdrop
Coordinator	✓	✓	✓	✓	
Medical device	✓	✓	✓	✓	
Network	✓		✓	✓	✓
Alarm system	✓	✓	✓	✓	