April 20, 2016

Docket No. FDA-2015-D-5105

Food and Drug Administration
5630 Fishers Lane
Rockville, MD 20852

Dr. Suzanne Schwartz,

Foremost in our response, we would like to applaud the FDA's efforts on cybersecurity. Over the last few years they have been a leading voice and a convening force to create a more resilient, safer care delivery ecosystem.

The latest medical advances lay at the intersection of patient care and connected technology. Integration of new technology enables innovations that improve patient outcomes, reduce cost of care delivery, and advance medical research. New technology introduces new classes of accidents and adversaries that must be anticipated and addressed proactively. The once distinct worlds of patient safety and cyber security have collided. Where the consequences of failure are measured in human life, we must all work together to know - not just hope - that our medical devices are worthy of the trust we place in them.

I Am The Cavalry has collected and curated responses to the FDA's draft postmarket guidance. Overall we feel the direction, tone, and tenor are appropriately tailored for both the gravity of the situation and the audiences the guidance speaks to. We hope to bring substantive new ideas to the conversation, from the perspective of the security research community, polished and shaped by conversations with many other stakeholder groups. We also offer observations for improvements or considerations that we feel could greatly improve clarity, adoption, and strength of the guidance.

I Am The Cavalry and the FDA share the common goal of uniting stakeholders to solve these difficult problems ahead. The FDA has led the way in outreach to the security research community, and I Am The Cavalry is encouraging and supporting ambassadorship from our community in reciprocation. We continue to pledge our commitment to fostering a high-trust, high-collaboration relationship.

Submitted respectfully,

I Am The Cavalry

Members of the security research community.
safer | sooner | together

Introduction

I Am The Cavalry is a global grassroots organization focused on issues where cybersecurity intersects public safety and human life. Our message is that dependance on computer technology is increasing faster than our ability to safeguard ourselves. Our areas of focus are medical devices, automobiles, home electronics and public infrastructure. Our mission is to ensure technologies with the potential to impact public safety and human life are worthy of our trust.

I Am The Cavalry is an initiative born from the cybersecurity research community, with participation from many different industries. Since our formation in August 2013, we have actively engaged the healthcare stakeholder communities in many ways. We held the first CyberMedRx[1] Multi-Stakeholder Summit on Medical CyberSafety in December 2015, participated in the FDA Public Workshop[2] in January 2016, also in January we published a Hippocratic Oath for Connected Medical Devices.[3]

Comments on the draft guidance

There is a lot to like in the draft guidance. Directionally, it matches successful practices in the software industry, in other manufacturing industries (such as automotive). Within the past few years, many in the healthcare industry have also been experimenting with these practices - developing them more fully and adjusting for their unique needs.

Coordinated vulnerability disclosure is part of the requirements for an incentive the FDA is offering to manufacturers. Public programs published by Philips,[4] GE,[5] and Draeger[6] demonstrate their leadership in the medical device industry. We have seen healthcare providers, such as the Mayo Clinic, engage security researchers to look for flaws in devices and make them known to the manufacturers. The Department of Homeland Security's ICS-CERT[7] has been instrumental in coordinating disclosures between researchers and affected manufacturers, and notifying affected organizations. The FDA published Safety Communications in May 2015[8] and July 2015[9] based on a demonstrated pathway to harm due

---

[1] CyberMedRx website https://cybermedrx.org/

[2] Public Workshop - Moving Forward: Collaborative Approaches to Medical Device Cybersecurity, January 20-21, 2016 http://www.fda.gov/MedicalDevices/NewsEvents/WorkshopsConferences/ucm474752.htm

[3] I Am The Cavalry Hippocratic Oath for Connected Medical Devices https://iamthecavalry.org/oath

[4] Philips responsible disclosure statement http://www.philips.com/a-w/security/responsible-disclosure-statement.html

[5] Security Concern & Compliance Reporting http://www.ge.com/security

[6] Draeger Coordinated Disclosure Statement http://static.draeger.com/security/

[7] The Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) website https://ics-cert.us-cert.gov/

[8] Vulnerabilities of Hospira LifeCare PCA3 and PCA5 Infusion Pump Systems: FDA Safety Communication http://www.fda.gov/MedicalDevices/Safety/AlertsandNotices/ucm446809.htm

to cybersecurity vulnerabilities. We are pleased to see this topic featured in the guidance and hope it can accelerate the industry trend.

We are pleased to see the FDA strike a balance on safety, privacy, and security. Not all vulnerabilities represent a risk to patient safety or privacy. By using a rubric of patient safety, the FDA allows for prioritized approach to addressing software flaws that allows manufacturers to evaluate these risks in the broader context of risks faced in delivering care, on a case-by-case basis. However, we do have some concerns worthy of noting, as well as suggestions for improvement.

Limited Accessibility of Devices to Vulnerability Finders

It should be noted that **vulnerability disclosure is heavily weighted toward accessibility, not criticality**. Vulnerability reporters are only able to identify flaws in devices they have access to. For security researchers, this disproportionately limits devices they can work on to those they can themselves acquire (ie. insulin pumps, not CT scanners). In healthcare environments, IT and IT security staff rarely interact with medical devices on a technical level, reducing their ability to proactively identify flaws or extant problems (such as malware on devices). The number and quality of disclosed vulnerabilities is likely to correlate closely to device cost rather than design security. A high number of low quality reports increases costs and timeline to assess and address patient safety issues.

Contracted vulnerability research may shrink this availability gap, at the cost of a restricted scope of testing methods. There is no single way to find vulnerabilities, and no researcher or project can identify all of them. Each researcher has different methods which yields different classes of issues. No single researcher, and no one project will find all flaws or even all critical safety flaws. Even to attempt an exhaustive search on a paid contract would be time and cost prohibitive.

> *We suggest the FDA - formally or informally - **encourage more open access to devices** by security researchers, in a way that does not put patient safety at risk. Existing practices and resources, such as cyber ranges, manufacturer "hackathons" on device designs or prototypes, healthcare provider security assessments, leveraging the Underwriters Laboratory, and engaging NH-ISAC can help address this access gap and are not mutually exclusive. These methods will not only facilitate discovering potential safety issues, they will broker stronger relationships between researchers and other vulnerability finders, with the broader healthcare ecosystem. The FDA may wish to convene a task force to look at ways to shrink this availability gap.*

---

[9] Cybersecurity Vulnerabilities of Hospira Symbiq Infusion System: FDA Safety Communication http://www.fda.gov/MedicalDevices/Safety/AlertsandNotices/ucm456815.htm

ISAOs are Not Yet Defined

The stipulation to **participate in Information Sharing and Analysis Organizations (ISAOs) is a hard requirement with a soft definition**. ISAOs are nascent structures, though they pull from existing concepts (such as ISACs). Their operational roles, goals, effectiveness, and timelines are undefined, as of yet. Their inclusion as a hard requirement makes clarity of purpose, outcome, and role a critical part of achieving the goals of the postmarket guidance.

> *We suggest the FDA **clarify expected outcomes from participation in ISAOs**. Some expected outcomes or activities may include the following:*
> - *Share vulnerabilities among organizations about commonly shared components.*
> - *Share practices among organizations to eliminate sources of vulnerabilities (for instance Common Weakness Enumeration - CWE) in design and implementation processes. For instance, test methodologies, tools, and resources for manufacturers or researchers.*
> - *Share information about software component testing results, robustness, etc.*
> - *Track incidents of reported vulnerabilities against addressed vulnerabilities for reporting. For instance, Mean Time to Remediate, or Mean Time to Patch, which can be compared across the industry.*
> - *Provide a "dead drop" for researchers to report issues anonymously when they feel manufacturers are intransigent or unreachable. This is a mechanism that will permit a release valve for critical issues in lieu of public disclosure when researchers may otherwise take that approach.*

Timeframe Uncertainty for Addressing Risks

We believe that **30 days is both too much and too little time** to address uncontrolled risk. This time frame is faster than Google's Project Zero[10] sets forward for vulnerability remediation, which is already considered both too short and too long. On the one hand, 30 days may be too short for proper testing of a technical fix or a workaround. On the other hand, 30 days may be too long when patient safety is at risk. A hard time limit may force one case or the other to be true, yet no hard time limit has no force for urgency when needed.

This is an appropriate point to mention the problem of "forever day" vulnerabilities, which are bound to exist until retirement of the device. While the guidance implicitly assumes that there will be an acceptable method to reduce risk to an acceptable level with the device in use, this may not always be the case. There may be situations where the essential clinical performance is tied to a capability designed in a way that inherently carries a high degree of risk, or the particular software component may not be updateable.

---

[10] Google's Project Zero attempts to identify vulnerabilities in Internet systems, and currently has a 60-day remediation window. https://en.wikipedia.org/wiki/Project_Zero_(Google)

In cases of forever day vulnerabilities, the medical device maker and the FDA have other remediative actions. For instance, recalls, product buybacks, and other actions can take an unsafe device out of use. We anticipate that these extraordinary situations will be the exception, rather than the rule. However, it is worth monitoring this situation.

> *We suggest the FDA **consider a multi-tier, multi-phase implementation**. We foresee several different situations where this tiering or phasing may be beneficial.*
>
> - *There may be threat level thresholds for a given vulnerability - private disclosure vs. public disclosure vs. known exploit in the wild.*
> - *Multi-stage implementation of a remediation - short term workaround vs. elimination of the software flaw.*
> - *A multi-tier approach may set different requirements for devices already approved for market, those in development, and those in design.*
> - *There may be designations for current state vs. desired state.*
> - *A shorter, or shortening, deadline may incentivize device makers to adopt and adapt modern device design and development methods. For instance, Agile and DevOps accommodate faster designs and fixes, as well as reduce burden and accelerate timelines for QA testing without sacrificing quality.*

Too Much Flexibility in Addressing Uncontrolled Risk

In general we favor a flexible approach to addressing uncontrolled risks. Writing this in April 2016, it is *impossible* to anticipate situations which may arise into the future, and *impossible* to build a one-size-fits-all approach to addressing security vulnerabilities. However, we recognize that flexibility may lead to two problems on the extremes - too much wiggle room to push responsibility onto providers; too much rope that prevents them from taking action or risks taking the wrong action.

All medical interventions introduce new risks in the course of addressing existing ones. There will always be uncertainty in modifying a device or the patient care workflow. In theory, medical device makers are in the best position to fully appreciate the impact of modifications and do so in a sound manner. The FDA has rightly placed the responsibility of making these decisions on device makers, and given them an incentive to do so. If manufacturers choose not to avail themselves of this mechanism, existing approaches (including recall) are still available.

More worrying to us is that device makers' remediation may leave providers in limbo between two untenable positions - an unworkable fix or an uncontrolled risk. Certain approaches to addressing risks may lead to high cost, heavy workload, or hard adaptations by healthcare providers or patients. They must then weigh two highly undesirable outcomes. When a significant burden is laid on providers and patients, can a risk really be said to be controlled?

> *We therefore suggest recognition of **implementation cost as a factor in assessing effectiveness of approaches to controlling risk**. When making decisions about*

controlled or uncontrolled risks, the primary concern is supposed to be patient safety. Similarly, a principle of least burdensome approach could estimate costs to a rough degree of accuracy as a component of patient safety - after all, each dollar spent on addressing uncontrolled risks is a dollar not available for care of patients.

Requirements Versus Guidance

There have been several discussions among security researchers, manufacturers, healthcare providers, and others about the need for the language in the guidance be given greater force as requirements. I Am The Cavalry does not feel strongly enough to take one side or another, but we feel compelled to introduce discussion points we have not heard elsewhere.

We have heard anecdotally that some larger manufacturers are treating the guidance as if it is a requirement, in an effort to be highly risk averse. Some are building capabilities to go farther because there is no penalty for failure at this stage.

Smaller manufacturers may not know about the requirements, nor have the internal capabilities to manage such processes right now. This may put their business at undue risk, not to mention patient safety implications of trying to rush something risky.

> *The current guidance status does not preclude elevation to requirements at a later date. In fact, **a multi-staged implementation of guidance can make requirements stronger and improve uptake**. This parallel experimentation approach without regulatory penalty for non-compliance may allow or facilitate:*
>
> - *Understanding of the benefit and risk of these practices on patient safety, manufacturer operations, healthcare providers, and other outcomes.*
> - *Leadership to emerge among medical device makers, establishing the art of the possible which may lead to a standard of care that exceeds the guidance.*
> - *Practices that favor for a prompt, agile, and secure response to software defects.*
> - *Design practices that seek to reduce known and unknown software defects before devices are developed and marketed.*
> - *Definition of poorly understood spaces such as appropriate timelines, ISAOs, and other softer requirements*

Further Thoughts and Suggestions

Point to Existing Resources

Published policies and programs in the software industry can serve as effective examples for medical device manufacturers. Use of vetted standards (such as those in the FDA Consensus Standards, including ISO 30111 and ISO 29147) and practices accelerate an organization’s maturity and ensure predictable, normalized interfaces to those who report

issues. Similarly, maturity models, such as the one published by Katie Moussouris while at HackerOne,[11] can help manufacturers understand a way forward.

A Software Bill of Materials

We are increasingly convinced that **a software bill of materials is critical to a healthy medical device ecosystem**. This is a list of third-party and open source software components used in the firmware and software of a device. This transparency unlocks free market forces, reduces operational burden, and extends device lifetime.

A software bill of materials unlocks free market forces at procurement, allowing patients and providers to make more informed buying decisions. Patients and physicians are in the best position to make decisions about course of treatment; impediments to this can only harm public health. Healthcare providers will be able to factor full costs or risks into procurement choices. Philips already provides a bill of materials to customers along with its MDS2[12] forms. And the Mayo Clinic reviews bills of materials as a part of their device procurement process.

A software bill of materials reduces operational burden and extends safe, useful lifetime of devices. The mean time to identify (and therefore address) safety-impacting software flaws is improved by facilitating comparison against common software components. This is true for brand new devices, as well as those past the manufacturer's stated lifetime, or indeed even if the manufacturer has gone out of business.

The FDA, manufacturer, and healthcare provider can quickly and easily answer the questions a) am I affected; b) where am I affected. The current spate of ransomware against hospitals takes advantage of a 9-year old flaw in the common software library JBoss. Yet there is no easy way to determine which healthcare providers are vulnerable and exposed. This significantly reduces providers' ability to address potential issues, and the FDA's capability to perform their oversight responsibility.

A software bill of materials renders risk and cost visible in a way that empowers a market to insure against residual risk. The Underwriters Lab Cybersecurity Assurance Program[13] includes this requirement. Likewise, the Financial Services Sector Coordinating Council[14] lists it as a requirement for cyber insurance.

A trusted third party can serve as an authority and reference for bills of materials. Subsequent analysis and reporting can unlock the benefits mentioned above without revealing all of the

---

[11] A Maturity Model for Vulnerability Coordination
https://hackerone.com/blog/vulnerability-coordination-maturity-model

[12] Michael McNeal mentioned this during the January 2016 FDA Workshop.

[13] UL 2900
http://industries.ul.com/software-and-security/product-security-services/product-testing-and-validation

[14] Appendix A
https://www.fsscc.org/files/galleries/FSSCC_Cyber_Insurance_Purchasers_Guide_FINAL-TLP_White.pdf

details publicly, and without superseding existing practices (such as that of Mayo Clinic and Philips). The FDA, NH-ISAC, or forthcoming ISAO may represent such a trusted party to act as custodian of current and historical bills of materials.

Conclusion

I Am The Cavalry applauds the work the FDA has done, and we hope it continues. Their leadership role from their position in the ecosystem helps make medical devices safer, protecting against cybersecurity threats to patient care. The draft guidance itself is a solid foundation to refine and finalize.

We renew our commitment to offer our to engage in a high-trust, high collaboration relationship with the healthcare ecosystem. Our mission is *to ensure technologies with the potential to impact public safety and human life are worthy of our trust*. And we are always happy to lend a hand to others who share our commitment to being safer, sooner, together.