UNITED STATES OF AMERICA

DEPARTMENT OF HEALTH AND HUMAN SERVICES

FOOD AND DRUG ADMINISTRATION

+ + +

CENTER FOR DEVICES AND RADIOLOGICAL HEALTH

+ + +

PUBLIC WORKSHOP - CONTENT OF PREMARKET SUBMISSIONS FOR MANAGEMENT OF CYBERSECURITY IN MEDICAL DEVICES

+ + +

January 29, 2019
9:00 a.m.

FDA White Oak Campus
10903 New Hampshire Avenue
Building 31, Room 1503 (the Great Room)
Silver Spring, MD 20993

FDA:

SCOTT GOTTLIEB, M.D.
Commissioner

SUZANNE B. SCHWARTZ, M.D., M.B.A.
Associate Director, Science and Strategic Partnerships
CDRH

AFTIN ROSS, M.S.E., Ph.D.
Senior Project Manager/Senior Science Health Advisor
Emergency Preparedness/Operations and Medical Countermeasures (EMCM)
CDRH

REID W. D'AMICO, Ph.D.
AIMBE Scholar
Office of the Center Director

This transcript has not been edited or corrected, but appears as received from the commercial transcribing service. Accordingly, the Food and Drug Administration makes no representation as to its accuracy.

Free State Reporting, Inc.
1378 Cape St. Claire Road
Annapolis, MD 21409
(410) 974-0947

PARTICIPANTS:

JENNINGS R. ASKE, J.D., CISSP, CIPP/US
SVP, Chief Information Security Officer
New York-Presbyterian Hospital

CHRIS BITZA, GSLC, GSIF, GSEC
US Product Cybersecurity Leader
bioMérieux, Inc.

JOEL CARDELLA
Director, Product and Software Security
Thermo Fisher Scientific

SETH D. CARMODY, Ph.D.
Cybersecurity Program Manager
Office of the Center Director
FDA

JOSEPH CHAPMAN, M.S.
Principal Hardware Security Engineer
The MITRE Corporation

STEVE CHRISTEY COLEY
Principal InfoSec Engineer
Trust and Assurance Cyber Tech Department
The MITRE Corporation

JOSHUA CORMAN
Founder, I am the Cavalry
CSO for PTC

GENE DASERIC
ICU Medical

ANURA FERNANDO, M.S.E.
Chief Innovation Architect
Medical Systems Interoperability and Security
UL LLC

CHRISTOPH FISCHER, Dipl.-Ing.
Sr. Systems Engineer Modeling and Design
Co-Chair IEEE 11073 PHD Cybersecurity Working Group
Roche Diabetes Care GmbH

BRIAN J. FITZGERALD
Senior Technical Manager
FDA

ALLAN FRIEDMAN, Ph.D.
Director, Cybersecurity Initiatives
National Telecommunications and Information Administration (NTIA)
U.S. Department of Commerce

CHRISTOPHER GATES
Principal System Security Architect
Velentium

JULIAN M. GOLDMAN, M.D.
Medical Director, Partners Biomedical Engineering
Anesthesiologist
Director, Program on Medical Device Interoperability and Cybersecurity
Massachusetts General Hospital/Partners HealthCare System

BILL HAGESTAD
Senior Principal Cybersecurity Engineer
Medtronic

ROBERTA HANSON, CISM, CISSP, PMP
Director, Medical Device Cybersecurity
Abbott Laboratories

MATTHEW HAZELETT
Biomedical Engineer
Implantable Electrophysiology Devices Branch
Office of Device Evaluation
CDRH/FDA

ZACK HORNBERGER, CISSP
Director of Cybersecurity and Informatics
Medical Imaging and Technology Alliance, a division of NEMA

KEN HOYME, M.S.E.E.
Director, Product Security
Boston Scientific

JIM JACOBSON
Chief Product and Solution Security Officer
Siemens Healthineers

MICHELLE JUMP, M.S.
Vice President, Cyber Program Initiatives
Nova Leah, Ltd.

TARA LARSON, CISSP, HCISSP, CSSLP, CISM, CEH
Senior Principal Systems Engineer
Medtronic

KEVIN McDONALD, B.S.N., MEPD, CISSP
Director of Clinical Information Security
Mayo Clinic

MICHAEL McNEIL
Global Product Security and Services Officer
Royal Philips

BEN MILLER, CISSP, GIAC, GREM
Vice President Threat Operations
Dragos, Inc.

COLIN MORGAN, CISSP, CISM, GPEN
Director, Product Security and Services
Global Product Security and Services Program
Johnson & Johnson

CARL NITCHY
Canon Medical Systems

LINDA RICCI
Associate Director, Digital Health Programs
Office of Device Evaluation
CDRH/FDA

DANA-MEGAN ROSSI, J.D.
Associate Director, Product Security Policy, Strategy and Incident Response
Becton, Dickinson and Company

ZACH ROTHSTEIN, J.D.
Vice President, Technology and Regulatory Affairs
AdvaMed

SABYASACHI ROY, Ph.D.
Director, Regulatory Affairs
BrainScope

ALAIN SILK, Ph.D.
Acting Diabetes Diagnostic Devices Branch Chief
CDRH/FDA

ROB SUÁREZ
Director, Product Security
Becton, Dickinson and Company

JASON TUGMAN, CISSP, CRISC, CCSK, ITPM
Vice President, Cyber Risk Engineering
Axio Global

EUGENE VASSERMAN, Ph.D.
Associate Professor of Computer Science
Kansas State University

CHARLES WILSON
Senior Architect
Draeger Medical Systems

BEAU WOODS
Cyber Safety Advocate
I Am The Cavalry

KEN ZALEVSKY
Head, Medical Device CyberSecurity
Bayer

INDEX

M E E T I N G

(9:02 a.m.)

DR. SCHWARTZ: Good morning. My name is Suzanne Schwartz, and on behalf of FDA, I want to welcome all of you, those who are in the room as well as those who are watching via our webcast to our fourth public workshop on medical device cybersecurity.

If I was asked to define discrete measures of success by which to describe the state of medical device cybersecurity in healthcare today, January 2019, I'd have to admit it would be a struggle. Now, don't get me wrong, that's not intended to be a critique of the state of the ecosystem, so let me explain.

Applying a set of quantitative metrics here simply wouldn't do justice to the evolution that we have all observed over the past 5 years. On the one hand, it falls short of characterizing the progress resulting from the efforts of so many of you across the broad stakeholder community on many different dimensions. And at the same time, quantitative metrics do not lend themselves to adequately reflect the intractable challenge areas that have kept us tethered from advancing and that held us back from true transformational change nor do they provide a window into new challenges that have emerged. This is, after all, like peeling back an onion and exposing layer after layer of complexity and recognizing that each layer has its own interdependencies and its own interrelationships that have to be considered as we look to build solutions together.

So this morning, to frame our current state, and as we look ahead towards the future, I'll invite you to accompany me as we take a brief walk down memory lane, and let's qualitatively assess where we are based on the review of the challenges that were voiced in our very first public workshop on medical device cybersecurity in October of 2014.

We split these challenges into two general categories: systemic, meaning that they tend to be universal concerns across all of critical infrastructure, and then stakeholder challenges, meaning that these were pain points perceived as unique to the healthcare and public health sector. Do these same challenges still plague us, or have we already overcome them? Have new, more complex challenges surfaced, or has recognition of prior concerns become that much more urgent for us to deal with today?

Arguably, we've seen incremental changes on multiple fronts, which has moved the ecosystem to a more secure posture while other areas remain works in progress. Yet certain challenges continue to hold us back, and they render healthcare a soft target for exploit of vulnerabilities, vulnerabilities that can impact on patient safety, on continuity of clinical operations, and on functioning of hospital systems at large.

At FDA we often hear from healthcare delivery organizations that while efforts to bringing new devices to the market with security-by-design requirements are greatly appreciated, it's the legacy technology that reside on their networks running on obsolete operating systems that are either unpatchable or fragile when attempting to update, which is of most imminent concern. Indeed, the global WannaCry attack that occurred in 2017, affecting hospitals across the UK and, yes, even some healthcare systems in the United States, shined a spotlight on this issue.

Interestingly, the legacy device panel that we convened back in October 2014 included the following questions for consideration:

- Can legacy devices be curtailed or phased out?
- Is this a viable goal?
- What impediments exist?
- Who's responsible for patching?
- What happens when a legacy device is no longer supported?
- Who should be responsible for securely configuring the device? And
- How can this process be improved?

I would ask how relevant and timely are these questions for us today? And yet you're probably wondering if this workshop is intended to focus primarily on the new draft premarket guidance that was released in October, why the emphasis on frontloading this topic of legacy devices? Is this not a postmarket issue? I'd like to propose that we reframe our thinking because legacy devices will continue to be in our midst.

As new technologies come to the market, older equipment inevitably becomes outdated. But that older equipment, does it not still serve a purpose? Does it still perform its intended functions? Wasn't it purchased as capital investments with the presumption that it would have a long use life? It would seem unreasonable and infeasible to exchange out devices that are older for newer ones unless the legacy device is brittle, incapable of being updated and patched, nor are applying controls a viable option. Consequently, that legacy device presents an unacceptable risk with respect to cybersecurity and safety.

The question is, then, how do we move from a state of legacy device fragility or brittleness to one of resilience? How might we design devices to be inherently capable of adapting and updating through their intended use life?

Resilience of medical devices is a total product life cycle matter. It is not exclusively in the domain of the postmarket. We must be able to address the importance of resilience at the time new devices are being designed so that even when a cybersecurity incident occurs the device is capable of containing the impact, protecting critical functionality, and recovering capabilities or services that were impaired.

And while resilience is one foundational cornerstone of advancing medical device cybersecurity and safety, trustworthiness and transparency are the other two key principles upon which the new premarket guidance is based and which we will discuss at great length over the next 2 days. Let's recognize that as is the case for resilience, both concepts, trustworthiness and transparency, are themes that really transcend any one phase and they apply throughout the entire product life cycle.

We've learned a lot, we've gained immensely from working closely with the security researcher, also called the white hat hacker community in recent years. We've learned the importance, actually the necessity, to shift from a culture and mindset of implicit trust to one of thinking like the adversary when it comes to designing new devices to better protect against the potential for exploit. This means incorporating comprehensive threat modeling using a systems-level approach during design of new devices, and it means taking the information that is gleaned through threat modeling to inform security-by-design requirements that ultimately can yield trustworthy medical devices.

And, finally, the third principle, transparency. Now, this is by no means a new concept. We've already talked about its criticality in postmarket, primarily through the lens of vulnerability, information sharing, leveraging ISAOs and the healthcare ISAC as well as in promoting the adoption of coordinated vulnerability disclosure policies and processes. Yet, to be most effective in safeguarding patients and in protecting healthcare systems, transparency needs to be baked in from the very beginning.

The bill of materials which we've proposed in the premarket guidance as a cybersecurity bill of materials, the CBOM, enables healthcare delivery organizations to make informed procurement decisions. And perhaps most importantly, it gives the institution information that it needs, well in advance of an incident, to manage its networked assets and to prioritize risk mitigations as part of its preparedness responsibilities.

Medical device cybersecurity wouldn't be where it is today without the hard work and heavy lifting that has been accomplished through collaborations and through many partnering efforts across the public and private sector. Some of these extend well outside of the healthcare community, and at FDA, we are grateful for these opportunities to learn from, to share experiences and best practices with other agencies as well as with other industry verticals, some of whom you'll hear from over the next 2 days. Indeed, we have a full plate of activities and crosscutting initiatives that are in varying stages of development or progress that you'll get to hear a lot more about from their respective leaders. We encourage you to make the most out of the networking breaks throughout the next 2 days to find out more about how you can contribute, how you can be involved.

One of the most impactful and rewarding collaborations that we've established is with our "I Am The Cavalry" colleagues, who serve as ambassadors to the medical device security researcher community.

Last year, Seth Carmody, my senior project manager for medical device cybersecurity who does just a stellar job, and I were privileged to be on site at DEF CON's Biohacking Village and to observe its fledgling medical device hacking lab in action, organized by "I Am The Cavalry" members. We both came away with the same reaction. Just imagine the possibilities if this seed effort could be scaled up to include increased participation of medical device manufacturers, engagement with healthcare delivery organizations, and the presence of clinicians, all in a well-controlled setting.

So we are most pleased to announce today the launch of the 2019 "We Heart Hackers" medical device challenge which will take place in August at DEF CON's Biohacking Village medical device hacking lab. This event will build on the varied principles of trustworthiness, transparency, and resilience in medical device security, and it aims to foster relationship building across researchers, device makers, provider organizations, clinicians, and patients. We welcome manufacturers to participate and to drive momentum by committing to take the challenge and inviting their peers to do so as well, starting today and going on for the next 6 months. Information about the "We Heart Hackers" medical device hacking lab challenge can be found on the links on this slide. We'll also have a panel session tomorrow that will take a deeper dive in explaining how the event is being organized, what are the benefits of participating, and how security information will be appropriately and properly controlled. We'll also have handouts of the invitation letter to be made available as well.

As is the motto of the "I Am The Cavalry" grassroots effort, teaming up in this way, we, all of us, can make healthcare safer sooner together. And I see someone at the microphone.

MR. SUÁREZ: Hi. Good morning, everyone. My name is Rob Suárez, and I'm responsible for product security at BD, and I'd like to declare that this year BD will once again proudly embrace the security research community at this year's DEF CON Biohacking Village and the medical device hacking challenge. I call on my other manufacturers to also make the same declaration today. Thank you.

DR. SCHWARTZ: Thank you.

MS. HANSON: Hi, I'm Roberta Hanson from Abbott Laboratories. Abbott will proudly embrace the security research community in this year's DEF CON Biohacking Village, "We Heart Hackers," and we call upon our peers in the medical device manufacturer industry to do the same.

DR. SCHWARTZ: Thank you.

MR. HAGESTAD: Good morning, Dr. Schwartz. Bill Hagestad. Thank you for having us all. Medtronic clearly embraces the independent security research world and commits fully to an elevated sponsorship level and look forward to this year's DEF CON. Thank you very much, ma'am.

DR. SCHWARTZ: Thank you.

MR. McNEIL: Michael McNeil, Global Product Security and Services Officer for Royal Philips, and we also definitely support the challenge and plan for our second year in a row to be participating in the DEF CON Biohacking Village, and we encourage all other manufacturers to participate as well in the collaboration with the researchers. Thank you.

DR. SCHWARTZ: Thank you.

MR. CARDELLA: Hello. Joel Cardella. I am the Director of Product and Software Security at Thermo Fisher Scientific. We were definitely the new kid on the block last year. We will be back this year, bigger and better, and we welcome you to come and try to break our stuff.

(Laughter.)

DR. SCHWARTZ: Well, I appreciate and I applaud each of you who have already come forward representing your company to take the challenge. A few of you, as you've stated, have already participated in last year's inaugural or fledgling event, and we do look forward to hearing you share your insights on what that experience was like tomorrow, when we take that deep-dive panel.

And I want to thank everyone here for the opportunity to share our FDA perspective on the state of medical device cybersecurity. I'm very eager to hear the discussions that ensue in our panel sessions. Each of you who are assembled here today in the room as well as on the webcast have an important, an integral, role to play in advancing medical device security. We're humbled, and we're delighted that you are part of this diverse community.

At this time I'd like to turn the microphone over to Aftin Ross -- I see that she is walking up the hall of the room -- who's going to provide further meeting logistics and guidance on the run of the day. Thank you very much and welcome once more.

(Applause.)

DR. ROSS: Good morning, everyone. Thank you again for coming to our public workshop. You might have noticed that on your name badge you have a number, and that is because we want this to be a very interactive meeting, and we'll be actually having breakout sessions on various topics that are pertinent to medical device cybersecurity. So that number on your badge actually tells you which group you are a part of. A majority of the groups are actually in this room, so Groups 1 through 18 are going to be in this room. You might see on the side that there are some numberings and some sticky pads and notepads and things; that's where the breakouts are going to be.

You already have an assigned facilitator, so they know who they are. But we will be looking for members of the group, one or two people, to raise their hand to help to take notes and help to write down the different ideas that have come up as part of that breakout session discussion.

Please note that we do have some starter questions for those breakout groups, but if there are particular topics that come up in your group that the group decides they want to explore further, please feel free to do so. We definitely want this to be an open dialogue and open discussion.

There are a few of the breakout groups that are not in this room, and they are in the side room, so there are a couple rooms outside this way and there is a room outside this way. When we get to the time for the breakout, we will show the schematic, and you can also ask one of us and we'll help to direct you to where you need to go.

We have purposely made the groups diverse, so we know that there might be several representatives from organizations that might have come here today. We would ask that, if possible, you try not to cluster together because we very much wanted to have a good mix of clinicians, medical device manufacturers, technology folks, clinicians, all those types of people within the same group because we feel like having that diversity is going to help to get better discussion. So please, as much as possible, try to stay within your group.

We very much also want to make sure that people know -- I know some people were asking about the restrooms. They are outside the double doors back here and to the right;

there are restrooms there. Also, when you came in, for those who might have missed it, there were program books that were on this side, so please feel free to take one. It has the agenda and the sessions and the speakers and all of that there as well.

So that's most of what we wanted to say about the meeting logistics. I think, as Suzanne also might have -- or you all might have noticed the fact that we have the microphones here, so we will have some opportunity for some question and answer during the day for the various panels that we have. And so please feel free to go up to the microphone when it's that question and answer period to ask your question.

And we very much hope that this is an opportunity for you to collaborate with your peers. That has very much been our theme, is that this should be a spirit of collaboration, so please take the opportunity to do that today.

I think we're going to go ahead and we'll get ready to transition. I just need to actually talk to Suzanne about one thing, and we actually might have one more announcement for you.

(Pause.)

DR. AFTIN: Okay, so our apologies. So it is January, and it is winter on the East Coast, and so sometimes we do have weather-related instances, and so we have been made aware that they have -- OPM, our Office of Personnel Management, has issued the opportunity for a 2-hour early dismissal because of some of the inclement weather, so we are going to be adjusting our workshop program to account for some of that. So we wanted to make sure that we let everyone know now is that we will have a working lunch. So for the pre-orders for the lunches, we would suggest that you try to do that, and we will have folks, once they get their lunches, to come back here into the meeting room so that we can go along through the next panel session. We also are going to end up doing a combined breakout session. So we're supposed to have two separate breakouts today, and what we're going to do, so that we can make sure that we get both of the topics discussed, is we're going to do them together. So we are going to have to modify our workshop program a little bit, so thank you in advance for working with us and having some flexibility here. We do appreciate that.

The other thing that I failed to mention the first time was that if you are online, there will be an opportunity during the breakout to have some online engagement as well, and so we will have some facilitators for that. So please don't necessarily think because you're online that you have to go away during the breakout. There will be an opportunity for you to engage at that time, also.

Okay, so I think those are the main things that we wanted to make sure that we articulated. We will try to see if we can't clarify for you somehow on our screens you can make your own travel arrangements, what the order is going to be now that we know we're going to have to truncate somewhat today.

So, again, I appreciate your flexibility in advance and in working with us on trying to make sure we get the good content, the good dialogue. We're trying to also make sure that everyone gets back to their respective home or hotel safely.

DR. SCHWARTZ: So we have a special treat today in that we are expecting momentarily to hear from the FDA Commissioner, Scott Gottlieb. And while we wait for him to arrive, I wanted to give people an opportunity, if they have any questions or thoughts, comments you wanted to share in response to my opening remarks, we certainly can do that right now.

UNIDENTIFIED SPEAKER: I just have one comment in regards to your remarks that you were making up front.

DR. SCHWARTZ: Sure.

UNIDENTIFIED SPEAKER: You mentioned a lot about procurement and working in the hospital environment and things like that. More and more devices are starting to get home healthcare, OTC type of things. Can you talk about how that may be integrated some more into this, as opposed to just necessarily focusing on just that kind of hospital environment?

DR. SCHWARTZ: You know, you bring up a very important point because we really cannot focus exclusively on the traditional healthcare environment, the brick-and-mortar facility, the hospital, given that patients are walking around with implantable devices, patients have devices that are connected to them at home and at the bedside, and the space where devices now really extend to goes, again, well beyond your typical facility. These are opportunities and challenges for us to discuss further as we move along, and I would encourage you to bring some of your thoughts to the fore around how we can address some of those.

I would say, off the bat, that this is one of the reasons also why we think it's rather important to incorporate the clinician and the patient perspective in the discussions that we have going forward so that the perspectives and the insights and what's important for patients and their considerations around devices are taken into account as part of the overarching benefit-risk assessment, the benefit-risk calculus.

All right. Wow, okay. It is a tremendous honor to invite FDA Commissioner Gottlieb to offer welcoming remarks to this FDA public workshop. Dr. Gottlieb has been a very strong advocate, indeed, a champion for us in helping us advance our public health mission and in also helping be that voice to the outside with the work that we've been doing in medical device cybersecurity. We know he's extremely busy with a lot of meetings to attend to, so really without further ado, I'd like to welcome Dr. Gottlieb to the podium. So please give him a warm welcome. Thank you.

(Applause.)

DR. GOTTLIEB: Thanks a lot. It's good to be here today. I appreciate the opportunity to be here, and it's good to be back in a reopened FDA. I want to begin by recognizing FDA's medical device cybersecurity team, especially Dr. Shuren and Dr. Schwartz, for their leadership on this important issue, as well as their unwavering commitment to fulfilling FDA's public health mission with respect to a whole range of issues, including the ones that we're going to be discussing here today.

I also want to say a few words about my gratitude to the people at FDA who really stood fast and persevered through difficult circumstances over the last 5 weeks during the shutdown. I think this agency showed tremendous grit. We stood fast to our public health mission, stayed at our posts even under difficult circumstances, and made sure that we continued to secure the public safety and that nothing bad happened on our watch. And it's very good to be back at a fully opened FDA, and one of the first orders of business is going to be this workshop here today. So this marks a special occasion.

Ensuring the cybersecurity of connected medical devices is one of our most critical device safety challenges and priorities heading into 2019 and into the future. As we've seen in recent years, the threat of cybersecurity attacks is not theoretical, and the risk of patient harm, whether from a ransomware attack that interrupts a hospital's operations or a hack that compromises a patient's device, is a very real concern. But we know that solutions to these challenges are not straightforward. Like the technology itself, there is not a one-size-fits-all approach to addressing these issues. Today's medical devices operate in diverse and often complex environments, from the bustling hospital ICU to the physician's office to patients' bedsides. Devices also vary in complexity, and device users are a varied group, including clinicians, patients, and caregivers.

Against this backdrop, the FDA's approach to device cybersecurity must be multifaceted. We must consider the implications of compromised devices across their complete life cycle, and we must operate in an environment of shared responsibility, collaborating across government agencies with industry, with security researchers, with patients and with providers. In short, with all of you and everyone here today. I'm proud of the Agency's efforts and their achievements in these regards, and I'm delighted that we're going to help build on its momentum this year coming up.

As we move forward, our efforts will be guided by the principles of transparency, resiliency, and trustworthiness. Transparency is essential as a building block in cyber incident preparedness and response. That's why last fall we issued a revised draft of our premarket guidance including new recommendations for manufacturers, manufacturer transparency by means of a "cybersecurity bill of materials." This is essentially a list of the software and hardware components of a device that could be susceptible to vulnerabilities.

The concept was born out of the FDA's experience working with device users such as hospitals and provider groups. Particularly in response to the WannaCry ransomware attack in 2017, we realized that a major challenge to efficient and timely threat response was that device users simply didn't know what they had. By providing this bill of materials, manufacturers would deliver much needed transparency. A bill of materials will enable device users or owners, such as hospitals and health systems, to more efficiently evaluate their inventory, identify devices susceptible to cyber events, and prioritize risk mitigation accordingly.

I know that one of the panels today will be specifically focused on this concept, including the types of information and levels of detail that should be included in these bills of materials, and we look forward to that discussion and feedback from this group.

We're also guided by the principle of resilience and trustworthiness because the technology develops rapidly, and devices, once purchased, may have a long life cycle. Today, one of our most critical cybersecurity challenges lies in addressing the safety risks posed by legacy devices. Many older devices were not built with cybersecurity in mind, and they may use insecure software, hardware, or protocols, leaving them vulnerable to attack. But unlike our laptops and our smartphones, many devices cannot simply be swapped out for newer models. Many require significant capital investments or intended for years if not decades of use and service. And although the hardware may be durable, many older generation devices were not designed with the ability to receive timely cybersecurity updates, including fixes and patches. They are, quite simply, not resilient or trustworthy.

That's a lesson we're committed to carrying forward as we develop approaches to ensuring device cybersecurity in the future. It's why we made resiliency and trustworthiness central themes of the updated draft premarket guidance. And that's why this August the FDA plans to participate again in DEF CON's "We Heart Hacker" challenge, a white hat hacker event. We've got to work on the names of these things next year.

(Laughter.)

DR. GOTTLIEB: I know that Dr. Schwartz has already discussed the details of this year's event, and I want to add my voice in encouraging manufacturers to demonstrate their commitment to the principles of device resiliency and trustworthiness by volunteering to take the challenge and participate with their peers at this year's event. And I want to recognize the extraordinary value that white hat hackers bring to the medical device ecosystem through efforts such as this one.

I also want to highlight the Healthcare Sector Coordinating Council's joint security plan, which was released yesterday. The plan articulates a consensus-based set of best practices for cybersecurity technology solutions for devices in a healthcare environment, providing a framework for manufacturers to assess their cybersecurity. In addition, it provides recommendations for hospitals to improve the security of their operations when purchasing and deploying medical devices on their networks. Developed by a partnership of government, industry, and provider groups, the plan exemplifies what the FDA and others can achieve through effective collaboration around these issues.

And, finally, on the concept of collaboration, as the FDA continues to expand our cybersecurity program and stay abreast of the evolving threat landscape, an important element will be engagement with clinicians and patients. A comprehensive approach to risk-based regulation must account for their perspectives and their preferences. For example, we must expand our collective efforts to increase clinicians' awareness and understanding of potential cyber risks, underscoring how device performance may be affected and, in turn, affect patient safety.

In addition, we recognize that patients live with their medical conditions, and they must make choices regarding their personal care and, as a result, patients provide a unique perspective which, to date, has been underrepresented in discussions among scientists and developers and regulators on how to address these device cybersecurity risks. To that end, I'm pleased that one of our panels today will focus exclusively on patient perspectives, and we'll continue to evaluate ways of eliciting and incorporating patient preference information into our regulator approach to device cybersecurity.

I want to thank you again for your participation here in this important workshop and I want to welcome you back to a fully operational FDA, one that is at full strength and will continue to be very strong today and into the future. And I hope you have a good discussion today and I look forward to your feedback. Thanks a lot.

(Applause.)

DR. D'AMICO: All right, everyone, we're going to take a quick break. If you could please be back to your seats by 9:50, we will continue with programming.

(Off the record at 9:37 a.m.)

(On the record at 9:50 a.m.)

DR. D'AMICO: Hi, everyone. Welcome back. I'd like to invite to the front of the room our Session I panelists and moderator.

(Pause.)

MR. HORNBERGER: All right. Thank you, everybody, for being here at the legacy device panel, Legacy Learnings: Drag of the Past Driving Increased Resilience in the Future. I was just informed that we are going to have a couple minutes cut off of the end of our discussion here based on the weather. So not too many, 5 minutes, and we'll just go ahead and roll in right away.

I think the first order of business is to have everybody introduce themselves. I will start. My name is Zack Hornberger. I'm Director of Cybersecurity over at the Medical Imaging and Technology Alliance.

MR. ASKE: I'm Jennings Aske. I'm the Chief Information Security Officer for New York-Presbyterian Hospital.

MR. McNEIL: Michael McNeil, the Global Product Security and Services Officer for Royal Philips.

MR. McDONALD: I'm Kevin McDonald. I'm a Director of Clinical Information Security at Mayo Clinic.

DR. SILK: I'm Alain Silk. I'm the Acting Branch Chief for Diabetes Diagnostic Devices at CDRH here at FDA.

DR. CARMODY: Seth Carmody, Cybersecurity Program Manager at CDRH.

DR. GOLDMAN: Good morning. I'm Julian Goldman, a physician at Mass General Hospital and the Medical Director of Biomedical Engineering for the Partners HealthCare System, and I lead a research lab on medical device interoperability and cybersecurity.

MR. JACOBSON: Jim Jacobson, Chief Product and Solution Security Officer for Siemens Healthineers.

MR. WOODS: Beau Woods. I'm with I Am The Cavalry. I'm also a cyber safety innovation fellow at the Atlantic Council, and I run the Biohacking Village device lab.

MR. MORGAN: Colin Morgan, Director of Product Security and Services at Johnson & Johnson.

MR. HORNBERGER: Excellent. Thanks, everybody. To kick things off here, just talking about legacy devices in particular, I know that it can be a contentious topic, and for the first couple of minutes I think it might benefit both the panel and the attendees just to talk a little bit about what we mean by legacy devices.

It was interesting here in these first two -- the first two talks delivered. Dr. Schwartz and Dr. Gottlieb both kind of touched on what a legacy device is or is not and some of the issues that surround them. But what I caught was that it seems the direction it's leading is that a legacy device, from that perspective, is one that cannot be updated or patched through an update mechanism. That's very truncated and perhaps a little bit simplified, but is that a workable definition from the panel's perspective or do folks have comments or opinions on sort of what legacy device is outside of that definition?

MR. McDONALD: Well, we're purchasing devices that can't be updated, so does that -- I mean, that then gives the intent that new devices that come to market are legacy devices.

MR. ASKE: What I would say is that, Zack, I tend to agree with your summary and, you know, some of the manufacturers who are here at this table are really forward thinking and thinking about, you know, partnering with HDOs on security. Unfortunately, we have manufacturers that come to us and still sell us or try to sell us, you know, devices that use Windows XP, and they sell them at significantly lower prices than the device running Windows 7, which itself is, you know, a little long in the tooth. So I think there's kind of a problem that many of the device manufacturers aren't getting in front of it the way that, you know, Siemens or Philips are in trying to be proactive. And so it's -- you know, I don't want this to be a pejorative term, I don't want to be seen as argumentative or being critical in some way, but it's just really there's a problem for healthcare organizations where the devices that are being brought to us by many manufacturers are using outmoded operating systems or they don't come with software update mechanisms. And so, you know, the FDA's guidance really pushing in this direction and the leadership we're seeing is really important.

MR. McDONALD: So we've just decided that if it is attached to our network, we just casually define it as a legacy device because when we look at it, it really has a lot of the other things that we need to do to it just to try to maintain it, the protections that we have to do, the monitoring that we have to do, so it's just an easier -- we just have the -- it's easier to be able to drop it in that bucket because it kind of acts, walks, and talks like any other legacy device most times.

MR. McNEIL: Yeah, and from a manufacturer's perspective, I would probably state, as Kevin just said, once it was put into your environment it immediately became legacy and it's how it needs to be managed and the life cycle of that device, I think is what's the critical discussion here. And, again, having a better proactive life cycle management approach with the device, within the architecture, within the -- you know, the hospital and/or the infrastructure that that device is being deployed, I think is what's the critical discussion points that we need to make sure that we have a good understanding of.

MR. MORGAN: Yeah, I would add that legacy means something different to everyone you speak to, every environment a device runs in, everyone who manufactures it. Even within our own company, legacy can have different meanings. Legacy can mean one thing for a device that we manufactured 15 years ago that's end of life and being moved away from, it could be a different meaning from a company we acquired and we now acquired their legacy security debt. So I think, really, it's hard to put a definitive term on the meaning of legacy. I think what legacy means is that you have, you know, to both Kevin and Michael's points here, is a conversation around what is the proactiveness that we're going to be doing around this? How are we going to maintain these new devices that we're launching so that the hospitals, the customers, can be confident that there are going to be patches made available and that, as a manufacturer, you can be confident that your customers are going to move off devices that can no longer be patched? So it's a broader conversation, and again, I think it really depends on, you know, who you're speaking to and how something is deployed or manufactured.

DR. CARMODY: Yeah, a lot of red lights here at this side. Yeah, I think it's really -- it's important to consider what it is and what it was, but I can tell you, in my opinion, what a legacy device shouldn't be and that shouldn't be a new device, for sure. And if I could sort of quantify what makes a new device immediately a legacy device is that it's un-updatable and what I mean by that is, in a more broader sense -- and I think, Michael, you were talking to it, Colin, you were just talking to it, is that for very good historical reasons, we haven't necessarily maintained devices for security and changing that, making it easy and fast not only for manufacturers to update devices in response to security risks, but also for hospitals as well. That's the key.

So if you can't change it, it's very difficult to respond to cybersecurity risks, if not impossible, and that sort of puts us in the sort of definitional ballpark of what a legacy device is and we can't be there in the future. So getting to Suzanne's -- like her slides from 2014, you know, let's start chipping away at some of these issues that I still see remaining up there and the first step, my opinion, opened up to the panel is updateability and patchability.

MR. McNEIL: So based upon that definition though, Seth, if I buy an extended support from Microsoft for an XP device that allows it to be updateable, then it's not legacy?

DR. CARMODY: Did they end extended support for XP?

MR. McNEIL: Not if you're willing to pay a king's ransom. It can still be negotiated, trust me.

DR. CARMODY: If it's being maintained for security --

MR. McNEIL: And let me make sure because, again, when we think about WannaCry, that extended support was gone, but they did issue a patch for WannaCry. So is that not -- how we would do with the definition there, because I think everybody would have put XP in that bucket.

DR. CARMODY: I think, in that case, which is a special case, I would admit that you're leaving to chance that Microsoft is going to come out with a patch for you, right? So I don't think that's a tenable situation, but I understand the use case.

DR. GOLDMAN: So I think it might be helpful -- first of all, yes to what Kevin said. Sure, we would buy a new device today or, as you've mentioned, Seth, out of the box it could be a legacy device because we're referring not about the medical device function, which might not be legacy, but specifically cybersecurity, of course. And so it could be a brand new device, it may be only the thing available on the market and sure, I think it's fair to say that it could be a legacy device because it can't -- it does not allow, it cannot be updated in some way to meet current cybersecurity standards.

I think that's the essence of whether something is legacy or not, can it meet current cybersecurity standards? And current depends upon the time that you're performing that assessment. You know, if for some reason in 2 years there were new threats that were not considered and if the devices can't meet those new threats because of technology limitations, then that would probably put them into the category of legacy. If the manufacturer can update the device or maybe alter the configuration, if it were a patch, upgrade, whatever it is, then it is able to meet the current cybersecurity standards. You know, an example is if you buy a brand new computer, out of the box it might be vulnerable before you patch it, before you update the software to meet the current vulnerabilities. It isn't the legacy device, even those it's vulnerable, because it's been designed to be updated when needed and to meet current cybersecurity standards. So I find that definition to be of some use, but there are gray areas, undoubtedly, as we pointed out.

MR. JACOBSON: So I would use a negative circular reference to define legacy devices as those devices that are not designed to manage some of the challenges I hope we will get into a little bit on this panel, and also that can be updated appropriately when that design doesn't meet that standard. So anything that's not in that space would be a legacy device.

MR. McDONALD: So while I don't want to, you know, sort of drop this in, but actually I don't care what we call it. This is a fascinating semantic discussion, but not one I want to waste time on. I would much rather drop devices into buckets of how -- what they have in their life -- how we're able manage their life cycle.

MR. HORNBERGER: No, and thank you, that's exactly where my understanding was coming to as well. It seems like, as we've talked here, the life cycle management has really come to the forefront as what -- and managing those devices that whether they're instantly upon procurement or at the end of the life cycle at a position where updating becomes an issue.

MR. McDONALD: Because these devices go everywhere from you can't touch it and do anything with it to it's fully upgradeable and will last 10 or 15 years and all of the in-betweens. So trying to put a single definition on it is not going to work.

MR. MORGAN: And I agree with you, Kevin. And to me, does legacy -- what does it even mean? I mean, does it mean life cycle management, does it mean your ability to manage the solutions for the longevity of them or is it really just about risk, because you may have a legacy device under a definition we might have used up here, but you could have mitigations in place that make the risks nonexistent. You could disable all USB ports, you could put a hardware firewall local in front of that box that is updateable and maintainable and now that medical device is, in essence, a bit isolated and are there risks? If there aren't any risks, does it matter if it's "legacy" or not?

DR. GOLDMAN: But it matters because of the additional effort to manage and maintain it and deal with all those issues. So I would separate the notion of legacy in something that is outdated to risk. Risk is --

MR. MORGAN: Exactly.

DR. GOLDMAN: Is it risk to the patient, risk to the network, risk to the hospital, risk to the PHI exposure? I think sometimes it's useful to lump them together, but often it's useful to tease that out for the reasons that you mentioned. But I still think a designation in a population of medical devices is really useful to know whether you have those that are kind of the special ones that need special attention, meaning a lot of extra effort to protect them and to protect the rest of the network.

MR. HORNBERGER: So I want to take, real quick -- go ahead and hop over. Is there a question based on --

DR. VASSERMAN: It's a question and a comment. I'm sorry for jumping the gun, I know it's not a question period yet. I'm Eugene Vasserman, Kansas State University. I ran up here when Seth defined legacy as not updateable and then the conversation sort of shifted in a direction I like more. I wanted to initially point out that some things are too updateable and the idea is to -- is it's supposed to be appropriately updateable, but then as I was standing here, I would like to pose a definition that's orthogonal to legacy or non-legacy and that is, is it software updateable? That is, is the hardware capable of supporting modern security standards whether that update is made over the air or manually via USB stick followed by the gluing of the USB port and so forth? So I might consider a legacy device one with open ports and it does not have enough resources to be updated to have a modern firewall or to have TLS communication capability. Whereas, a device that has enough resources is over-provisioned and is capable of supporting TLS but just does not support it at the moment, you could apply a patch and make it support it, I would consider that moderate to non-legacy. Maybe there are levels, as Kevin said. You didn't say levels of legacy, but I'm misquoting you on purpose.

MR. WOODS: So one of the things that I'm hearing, just from a little bit standing back and observing, is we're talking a lot about -- as Kevin rightly pointed out, you know, I really don't care what we call it, but there's got to be something. So I've heard us talk a lot about the idea of roles and responsibilities, whose job is it to do something and at what time, at what threshold. Capabilities, what are the -- what's the art of the possible with this device, whether it's brand new, off the shelf, or older. And then what are the adaptive -- what's the adaptability of the environment to support a device that doesn't have the native capabilities? I wonder if that's maybe not a better framing than a term that none of -- you know, we've got what, 12 people on the panel and we've come up with at least 15 definitions so that, in itself, I think, says we might not get this question answered by the end of this panel, especially abbreviated because of the weather.

MR. HORNBERGER: I think that that's an excellent point, and I don't want to belabor the definition of it any longer, but I do think that it makes the discussion difficult if we're -- if we do take a look at those two buckets, if we were to take the two buckets of software-updateable devices at any point in their life cycle or devices that are not software updateable at whatever point in their life cycle. I think those two things you mentioned, roles and responsibilities, capabilities, adaptability, I think that those three topics within those two buckets would lead in different directions.

So I guess my question to the panel at this point would be is that accurate, or do you think that those three topics are similar across both types of devices? You have the software -- well, the non-software updateable.

MR. ASKE: So, you know, in thinking about this presentation -- and I'll get to your question, Zack -- I was thinking about this really from the challenge of the healthcare delivery organization and the state of things, so to speak. So relatively recently, New York-Presbyterian acquired a regional hospital, community hospital, for another term, and you know, we looked at the risk posture of the institution and we could see that they had really done nothing about their medical device security debt, if you want to call it that. And without the acquisition from the hospital from Manhattan, they probably would not be able to do that, right?

And for folks who aren't as familiar about this, hospitals operate in pretty small margins and, you know, this was talked about earlier today but really, when hospitals are making these capital investments in devices, they just can't swap them out when there's a security vulnerability.

So, you know, I go back to the case of some infusion pumps that just about a year and a half ago were announced to basically have a security vulnerability and there was no software update mechanism and we were about to buy hundreds of those devices, and luckily, I saw the notice and reached out to our procurement folks and we stopped it.

So this concept of software updateability is really important and, you know -- and to beat the proverbial dead horse, yeah, I don't know what legacy is. But, ultimately, there's a problem here, right, where we have this healthcare delivery system in which medical devices are really important, they save people's lives. Security was not part of the initial kind of design focus for this, and it's shifting. You know, the manufacturers are really starting to step up, but there's still technology and security debt and, you know, ultimately it's kind of like how do we, as partners in the HDO side, in the manufacturer side, figure this problem out? And as best practices, you can see emerging from groups like, you know, the HISAT trying to advise how to deal with these legacy scenarios. We've heard some of the discussions from some of the comments today, but ultimately, there's a problem statement; it's more around roles and responsibilities and risk than it is about, you know, defining this neatly.

MR. WOODS: So maybe one of the ways to look at it is if you take it from that roles and responsibilities perspective, at what point does the role and responsibility shift to a point where it is unreasonable for the person in the current role to do it, right? Like, I can't write a Microsoft patch, as a hospital, no matter how hard I try. So it would be unreasonable for me to support that device if I have to write my own Microsoft patches, for instance.

On the other hand, you know, a device straight out of the box that has no update capability, maybe there are updates coming out that I could avail myself of but the only option is replacement of one or more components of the device because it's not software updateable.

So maybe framing it differently and pulling on the thread that Jennings brought up on those roles and responsibilities, that might let us get to a practical workable conversation and I'm not talking just about on this panel but more broadly, more wide scale. And it also might mean that organizations that are larger might be able to deal with certain legacy issues, legacy related issues that small organizations wouldn't be able to. So there wouldn't be a single definition of legacy.

DR. CARMODY: Beau, you bring a really good point up. I'd be remiss in not pointing people to the premarket guidance where we defined end of life and end of support. Actually, I think we may have gotten rid of one of the definitions in there, but generally, what is the thing that we expect manufacturers to do when it becomes untenable to support something if they're reliant on Microsoft to release a patch to make it supportable? And we need help there. So I hope the conversations do go to that today and then feedback to the docket on what -- when does that transition point happen, when does that risk transfer happen between the medical device manufacturer and the healthcare delivery organization?

DR. GOLDMAN: I think the hardware question is just worth noting. We have had devices that are -- that could be brought up to current security standards with a hardware upgrade. They're almost not the same device. They look the same on the outside, but the upgrade involves substantial increase in memory, for example, of adding new boards and flashing the memory. So I think there's value to recognize the difference between a software upgrade, software upgradeable to meet current standards and, you know, the need for hardware. It's kind of a different bucket. It might be too expensive, it might take too long, the resources might be quite different. So things to consider when adding the hardware point to the discussion.

MR. McDONALD: So from an HDO perspective, while again this is a fascinating semantic conversation, what are we trying to accomplish by dropping something into legacy or not legacy? I mean, really, if I look at it there's two things. You know, actually, there's one thing, what kind of risk does it have to me after I'm able to apply kind of compensating controls that I have and that -- whether I put it into a legacy bucket or not. So if we're looking at trying to define things for regulatory purposes, okay, let's have that discussion. But as far as us getting it in, if I put it on the network, I don't care whether I call it a legacy or whether I call it a new device. I look at that device, look at the individual attributes of it and try to figure out what we need to do to maintain it and what we need to do to make it secure for our patients.

DR. SILK: Kind of along those lines, I want to just get back to a comment, a word that Seth used, which is expectations. And I think it might be nice to hear from the health delivery organizations here, what their expectations are for the devices that they get and their upgradeability and their future security.

MR. ASKE: I don't have an easy answer to that because I think I'm very sympathetic to the fact that, you know, we can't ask for a device that we bought 15 years ago to support the latest revisions of Windows or what have you. You know, I think there's a real challenge here, and I want to touch on -- you know, Kevin asked like why are we discussing the term legacy and I think this is part of like a cultural narrative, right? This is very much akin to the discussions that we had around automobile safety and, you know, one point, it wasn't part of the design and manufacturers resisted and eventually people figured out we can make money by making our cars safer and it became a selling point, you know.

And I think the reality is that we're just -- this is part of a discussion, but there's a real problem which is there are medical devices at my hospital and, Kevin, at your hospital, that it's difficult to secure. We spend a lot of money and resources doing it and it's a challenge to patient safety. So I think that the value of legacy is to acknowledge this is an issue and that manufacturers need to adopt the principles that are in the guidance from the FDA around making security part of the design and that's -- you know, the work that we're doing which people will hear about later around SBOM and stuff. I mean, there's a lot of good work happening now to try and address this.

DR. GOLDMAN: I agree, and I think part of the benefit of defining legacy is scoping, helping us to define what we want for the next generation of devices. So even if we never talk about legacy again, we could change the conversation to what is needed in modern medical devices to support cybersecurity standards. Then we wouldn't use the word legacy, but we would be talking about the same thing.

MR. JACOBSON: And speaking about life cycle planning, Seth talked about the transfer of responsibility and how that is a difficult thing to define and to manage. The key to doing that, though, is not to make it a single point in time, right, to make it clear that you -- when the end of life of a product would be, to provide information ahead of time, say from the beginning, from purpose onwards, about the security posture of the device and what the HDO would expect to see over time and especially at that point of end of life.

MR. McNEIL: Yeah. And, again, I agree with Jim there in terms of the transferability discussion. I think we also, as manufacturers, need to do a much better job in that level of communication and working with the health delivery organizations around that expectation component. I know one of the key factors that we look at from, you know, an internal perspective as a manufacturer is to try to make sure that we can sync up between some of those expectations, because it becomes very tenuous if I am selling a device and potentially servicing a device that goes well beyond what that life cycle would be, and it sends and makes messages from a communications perspective and in expectations as it relates to some of that transferability.

So one of the key areas that we're trying to get in focus is to get better alignment. And for us, you know, to admit either I'm going to maintain and be able to support, you know, this particular solution, the ability to do the updates as a potential part of the definition and the description or I'm not, and make sure that, you know, we have the very frank discussion and points around that so that we can move the life cycle component and the management of it to be something that is understood and not some, you know, enigma that we're trying to gather information on.

MR. MORGAN: I would add to that. You know, it's not just our regulators like the FDA or those around the globe that are driving this; it is the customers, the hospitals sitting at this table, the risk assessments the contractors use. And we just heard Jennings say that they shut down, you know, a purchase because the device was not updateable. We're seeing it across the board in contracts, what version of operating system are you running, what's your plan to get off Windows 7, as part of that negotiation. So it's just the natural progression of our industry where you look at the timelines of how FDA has driven this in your 2014, the premarket/2016 postmarket. Now, 2019 premarket, again. Over time, we'll eventually start getting to a better state where newer devices are being rolled out with security capabilities. And frankly I always say, you know, premarket, that should be the easy part. Building security in, we've known how to do this for 10, 15 years just across the attack industry.

The postmarket is continuously a challenge, and we've got to figure out how to build sustainable models to maintain these devices in a fashion that (1) are updateable in a reasonable amount of time but are not done too frequently, that cause potential harm, (2) that are communicated appropriately to our customers so that they're aware of what's coming, and (3) ensure that across the spectrum of the healthcare industry there's some type of communication across that board so people are aware of what's happening.

And we're seeing trends head in that direction. We heard several manufacturers, you know, talk about the DEF CON Biohacking Village this morning and their commitment to it and we've seen the Healthcare Sector Coordinating Council publish, a few weeks ago, recommendations for hospitals; yesterday, our joint security plan for medical device manufacturers. So it's a continued -- the continued commitment from this industry to work together and figure out how we solve these problems is really kind of the guiding principles I look at to help drive this forward.

DR. GOLDMAN: Colin, to that I'll add the comment that we certainly do that within our healthcare system, we've been performing cybersecurity assessments for some time, and I think it's worth noting that manufacturers have become much responsive and much more thorough in their sharing of information now about strengths and weaknesses and their plans going forward. So there's been a positive --

MR. MORGAN: And it makes a difference for our businesses. They see that and they make changes. We have some, you know, some aspects of our business that are not getting hit yet by customer assessment questionnaires or contract reviews, and it's a little more difficult to get them to change versus those that are getting hammered with multiple assessments per week. They see the writing on the wall and, you know, it's just a matter of continuing that change.

MR. HORNBERGER: So transparency, communication, and expectation, I think, are three critical items that have come out of the conversation so far and I liked -- you know, I was kind of Frankensteining together a couple of the ideas pitched here. It sounds like what's needed is a modern -- what is needed in a modern medical device to support lifetime planning and risk assessment or cybersecurity might be the question to approach here as we get towards the end of the panel

We were talking a little bit about assessments, but are there other things from folks on the panel, you know, in response to that question of what is needed in modern medical devices to support those sorts of things, lifestyle -- I'm sorry, life cycle planning, risk assessments and trying to address some of those concerns?

DR. CARMODY: So I know it's a bit nuanced, but I'm going to bring it back to updateability. You'll hear me talk on the threat modeling panel and possibly in the market overview about -- you know, we talk about analysis. When designing devices up, what is the -- our analysis is by the troubles in security is incomplete, we just can't predict what a threat will do. Therefore, updateability is our most promising mitigation. If it's not updateable, how do we respond to risk? If it's not easily updateable, how do we respond to risk? So I think for me it always comes back to that.

MR. MORGAN: I would add just one little twist to that, too. And there was a question earlier when Suzanne was speaking in the beginning about not just focusing on devices in hospitals, but as we're seeing just technology infuse the commercial -- I mean, I'm a victim of it, and I've got every device you can imagine in my house from smart thermostats to smart cameras, and such and it's perpetuating a larger problem where I don't know when that device is going to be end of life, and what happens if it is, I got to get up on that ladder and go really high and it makes me uncomfortable and change out that motion sensor again.

But, you know, there are growing areas inside of our community that are building over-the-shelf-type smart solutions that have Bluetooth that connect to phones and these are and may be considered medical devices but they may be considered low risk. So we have to always make sure to factor in solutions like that and don't always consider that every device is extremely high risk or is always connected to a hospital network. There may be lower-risk devices, smaller units that are very cheap to purchase, and there may be decisions that the manufacturer has to make based on, you know, what that device is supposed to do. So I just wanted to ground that, you know, we always have to focus on the broad range of where the medical device sector is right now and where it's headed.

DR. GOLDMAN: And also consider the clinical operational impact that kind of builds on what you're saying, Colin, the clinical operational impact of even if something is software updateable there still could be a substantial impact. For example, perhaps the upgrade process takes 30 minutes and requires and breaks -- you know, temporarily interferes with connections of, say, bedside devices to monitoring systems, alarms and central stations, which requires a substantial effort on the part of a healthcare delivery organization to prepare and manage that. So that needs to be the part of the overall picture that manufacturers should consider. Just checking the box that it's software upgradeable but, you know, it isn't really feasible in a clinical environment would be an issue.

MR. JACOBSON: So talking about transparency is important because what's critical is that the information that a manufacturer will provide to an HDO can be integrated into the risk management, risk management system, the risk management profiles and processes that are maintained by the HDO. So it isn't a matter of just having a checkbox and saying yes, we're a trustworthy device, trust us. It's providing information that's actionable. CBOM is an example.

MR. McDONALD: You go, Jim. That's it, transparency, so we can make those risk decisions ourselves instead of allowing manufacturers to be able to make those risk decisions for us.

DR. SILK: I mean, another challenge here is that, you know, there are also high-risk devices in use by patients at home and these aren't just restricted to the hospital setting, and transparency can be a little bit harder to communicate when you're talking directly to a patient.

MR. HORNBERGER: That's definitely true, but I think that might be a topic for another panel to consider. We have a couple minutes left, so I want to make sure you have time for questions.

MR. CORMAN: All right. I don't know if this is on, but maybe it can be repeated. Josh Corman, founder of I Am The Calvary and a chief security officer in the supply chain. Security officers tend to talk casually about legacy is a risk decision you didn't make but inherited. So in casual speak, that could be an operational decision of the things we've inherited but didn't explicitly decide to do in this modern threat landscape and environment. But in a more prosaic way, up until 2014 we had zero FDA guidance on cybersecurity for premarket, and up until 2016 we had zero for postmarket. So, really, anything that was designed or developed prior to those was security by accident. And maybe more practical for HDOs and healthcare delivery is we should cleanly differentiate devices that were not ever designed in the era of guidance from FDA. Now, this guidance is rolling, and we're here today to improve the premarket guidance. We had the Mirai botnet show a hard code of passwords. We know software bill of materials were improved, the risk decisions of clinical environments, as to whom they purchase from and how long they maintain it.

In our taskforce, we specifically said the durable goods that are supposed to live 15-plus years will not have the same threat model as something like this, which is supposed to last two. So the software life cycle is wholly different than the physical durable goods, and we still need academic and grant challenge to come up with long-term lifetime support for medical devices.

So rather than quibbling over the definitions, I hope that we all look at devices that were -- I mean, the 2014 guidance isn't even in the market yet, right? Think about that. Sort of. You had a couple early adopters getting ahead of this, but there wasn't actual guidance, so we still have a couple more years before we see the first fruits of that. And I'd like to suggest that since software ages like milk, not like wine. We should take a more reluctant risk posture towards things that were pre-guidance and we should -- when we update the guidance, we should consider those behind current guidance for their lifetime.

MR. GATES: Chris Gates, Velentium. Besides being very male oriented, this panel is also very HDO oriented. This is the premarket guidance, not the postmarket guidance we're here to discuss today. So as a representative of dozens of different device manufacturers, I get the privilege of seeing a lot of how they do business, and also, being in the industry for over 40 years has given me a lot of insight into how business works. There are very few truly new products. In my business today, fortunately, I see a lot of those, but most of those are incremental improvements, and the premarket guidance doesn't once mention legacy, the October version. So how are we supposed to deal with this? Many of these things are in process currently, or they're slight modifications to existing products. In some cases these are implants. Are we suggesting we go force the patient back into surgery to pull the implant to make it stronger because an inductive communication is not authenticated or integrity checked? I don't think anybody believes that's a smart approach to it. Yet the new premarket guidance doesn't leave room for any of these. So instead of the HDO-centric approach, we really need a manufacturer-specific approach. And Michael is the only one sitting up there who brings that and I appreciate that, but we really need to think about how we handle those.

MR. WOODS: I think the device makers would disagree with that.

MR. GATES: Okay. Well, so far it's been talking about PC-centric and updating Windows execute. That's not a common thing that we encounter. Usually it's I've got an application-specific IC that's implanted in a body, how do I deal with this because it's 10 years old?

MR. McDONALD: So interesting comment because I've been having those discussions about implants over the last couple of weeks, and I think the thing that really throws us off in these discussions is taking the widely varied comments of, oh, yeah, you should be able to update everything, too. Yeah. Well, quick -- you know, telling us what we've got to do because they're implants and because it's an implant it's never, ever going to work.

Well, guess what? Sometimes you can make updates to implants. Sometimes you have to look at the risk based upon it, sometimes while you aren't obviously going to pull it out because there's risks involved with those as well. So you have to be able to, in a reasonable, rational way, take a look at the risks of doing an action of something versus not doing an action of something, apply all of those compensating controls. Discussions I've been having about some of those are the risk is greater if you pull the implant out than if you leave it in. But, oh, by the way, here's some process steps that we can do to be able to decrease that risk. And my experience with the FDA so far is that I'm pretty sure Suzanne's not going to tell, you know, all the pacemaker patients that they're going to have to be able to go and rip that stuff out. So it's being able to live with that shade of gray, is how we're going to be able to move forward effectively.

MR. MORGAN: From a manufacturer perspective -- and I know we've got times up, so I'll make it brief -- is, you know, we view the guidance whether it's from FDA, TGA, Health Canada, CFDA, South Korea, Japan, PMDA. It is, to us, the requirements whether you're legacy or not. So if you're a legacy device going through an incremental update, if you go and have to resubmit that, you're going to get asked cybersecurity questions. CFDA is going to push back on you, they're going to ask you what you did for your threat model, what did you do for your requirements. FDA is going to do the same thing.

So you can't look at it through the lens of, you know, incremental versus major change versus, you know, a brand new feature set, brand new products, you have to build these in no matter what. And if it's the legacy device, you have to figure out how to fit them in or you've got to figure out how that gets added to the next version and the next model. That's just the way it is. Hospitals are going to do the same thing. They're not going to care that this is incremental version 1.1. It's a brand new product to them and they're going to say why didn't you use 3.0? Oh, you built this before or you've got to add these things in now or you've got to do it by this date. You know, July 2020, if you don't put these features in, we're going to remove your product. And that's the reality of what's occurring in the industry right now.

MR. HORNBERGER: So it looks like time might be up. Do we have time for one more question, comment? No? Sorry, that's going to be it for the panel. Thank you, everyone, appreciate you coming up here and discussing this.

(Applause.)

DR. D'AMICO: And next up, I'd like to welcome Seth to the podium to discuss Medical Device Premarket Guidance Draft Overview.

DR. CARMODY: Folks, if you were displeased with my panel performance, I'm very sorry, I've been upgraded to a talk. So apologies in advance. All right. So we're here today to talk about the premarket guidance. I won't belabor it. I'm just going to give a high-level overview to introduce the concept, if you're not familiar, and maybe provide a bit of substance and color around where they came from. Next slide, please. Oh, I have a clicker, look at that. All right, onward.

You know, I think Josh mentioned that there's been this steady drumbeat of guidances coming out and policies coming out from the FDA and that's just the nature of cybersecurity and it has to continue to evolve. So like designs and risks evolve, so do policies. So we've taken basically everything that we've learned in the last, I would say, 5 or so years, even extending beyond 5 years, and tried to package it within this premarket guidance, and it comes from interactions from all, with all of you. It's been a collaborative process since the beginning, and it will continue to be a collaborative process. That's why we're here to discuss these things. The docket is open, please go there. After the discussions today, we need your feedback. We need to make this the most sound policy that we can, and we can't afford to miss the mark on any of the elements that are in here.

And I would like to thank you all, and then the members at FDA who have contributed to this, both in the review branches, cybersecurity working group, and all the hard work and thought process that goes into these documents, it's quite incredible to be part of, and there are some really hard questions and challenges within here, so we need your help. Please contribute.

So taking all of the things that we've learned, basically, we've packaged it in this document. And I had the privilege of meeting with a gentleman well versed in threat modeling who said you know what, I just realized your whole document is a threat model, and I'm like yeah, that makes a lot of sense. We've taken all of our experiences, and we said here are the threats to medical devices, and then we give you a list. We say, hey, here are the things that you should be doing to mitigate against those threats. It's risk based, which provides you a bit of flexibility in terms of what you're deploying. We can talk about that in a second. But it tells you what are the things that we've seen that are effective in dealing with the threats that we've seen to the medical device space. And that's really what the document is about. And if you don't think the document gets there, please let us know; we can't afford to miss the mark. So there are numerous examples, routine management stuff, worldwide attacks, we've tried to bake into this document.

When we started the journey to revise this document, there were still problems in the space, and I wanted the source of those problems to be the premarket document. I said maybe it's the premarket document that's causing all our troubles; maybe we missed the mark. And when we went back and we reevaluated it, it was really sound, and in fact, we kept literally all the content, maybe a few tweaks that you'll find in the 2014 document. It references the NIST five core functions. After all, we're trying to make medical devices more secure. When they're deployed into the hospital environment, they can be managed as endpoints. Hospitals leveraging a cybersecurity framework like NIST have to do the five core functions. We want medical devices to enable them to do that. So keep that alignment.

It has a similar structure and flow overall, and it's focused on documentation. Documentation is really a double-edged sword. It creates a marker in time of that process that you went through that demonstrates that you went through the actual processes that are outlined, and then again, it's laborious, it's hard to get documentation right, but you have to submit it to the FDA as part of your premarket notification. That's the only way that -- well, it's one of the primary tools we have of communicating to the FDA, so it's very important. And we talk about threat modeling later, like what does threat modeling documentation look like to the FDA? Have you ever submitted threat models to the FDA? If you haven't, you have a lot of conversations to have internally in your organization.

I mentioned that a lot of the concepts were sound and really just sort of needed expanding. I make it sound simple, it's difficult, but we really wanted to -- we wanted to collaborate and focus on the concepts. You might find CIA -- confidentiality, integrity, and availability -- but we wanted to add a piece that we thought was so critical and that was the authenticity. Not new to security, for sure, but definitely something that we felt like, as we've just mentioned in the legacy panel, what one of the things that was missing when devices weren't designed with security in mind was how are you saying -- how are you proving trustworthiness? Something's designed to be -- if something is talking to you on the medical device, your medical device, you're telling me to do something, why should I be doing that thing that you're telling me to do? What are the mechanisms in place that authenticate, right? Who is telling it to do that? Is there some authorization mechanism?

So designing trustworthy devices, that came from the security community. How do we get these designs to say, yeah, I should be listening to you, I'm going to offer a bolus command because I know that you're authenticated and you're supposed to be telling me that I should be delivering bolus commands.

The second big thing is about multi-patient attacks. One of the things that concerned us, as a public health agency, is the type of vulnerabilities that allow your devices and systems to impact many patients either simultaneously or nearly simultaneously. So, again, that's part of our threat model, and you'll see that in the Tier 1 controls that those are the things that we offer there, you know, a default asking you to do, are basically getting at that, they're trying to prevent or knock down the risk of multi-patient attacks. There's a tiering system, higher cybersecurity alert risk and lower cybersecurity risk. Just to make sure everybody's on the same page, I know this will be a topic of discussion today, thank you. This is by default, we want you to go through an analysis, a comprehensive systematic security analysis of your design, and when you do that with reference to a threat model, you'll sort of have the answer whether you're Tier 1 or Tier 2, and we'll sort of lead you to the answers. What am I doing, what am I building, what are the risks, and what am I doing about it?

Plenty of stuff to talk about on cybersecurity bill of materials, and I've heard that you guys love the name. That's all up for discussion and I don't want to get hung up on it, but you know, certainly we wanted the messaging to be right, but it's very important for transparency.

We'll talk about system-level threat models. I'm going to forego the nitty-gritty of the criteria. It's a threat model. It's basically saying are you connected and can you cause harm to multiple patients, direct harm to multiple patients. That's really what it's saying and by default you're there, so your designs have to start there. When you submit to the FDA with documentation, you have to provide that design documentation and the rationale of why you haven't done something that we've put in the guidance. We're also humble folks. We may have missed something. There may be something different to do. The latest and greatest security technology comes out, you should be able to rationalize and provide us that rationale of why you're doing that for design.

And then I can't miss -- I'll conclude with, finally, the pre-submission process is something we've been having people do and engage in. If you have questions, timing is so critical in this point, you have to engage the FDA early and often. Don't come to us at premarket time and say I got the design wrong, you don't want to come to that conclusion and have to go back. So the pre-submission, get ahead of it and let's have that discussion. Thank you.

Do I have time for questions? Are we cutting it? Is it the snow? It's the snow, yeah. They won't say it, but I know. Thank you.

(Applause.)

DR. D'AMICO: Thanks again, Seth. And next up, I'd like to welcome the panelists and the moderator for the Session II Plenary Panel to discuss Threat Modeling and System Approaches.

DR. CARMODY: Could I have the threat modeling panelists come up, please? Can I stand? I'll stand. Yeah, please. Thank you. Ladies and gentlemen, if you didn't like my presentation, I regret to inform you that I am now the moderator for the next panel.

(Laughter.)

DR. CARMODY: Great, thank you. Excellent, we have all our panelists, yes. Let's get started. For modeling and systems approaches, this is the title of our panel today, I am your most excellent moderator, I won't be introducing our guests, but I will let them introduce themselves. I will do a just slight framing. I talked about threat modeling. In reference to the guidance, it is a threat model; it's a way to systematically assess risk threats against your system. I think I'll let the esteemed panelists talk about it, and I'll turn it over to them to introduce themselves, and then we'll get started.

MR. SUÁREZ: Hi. Good morning, my name is Rob Suárez, and I am responsible for product security at BD, Becton Dickinson.

MR. TUGMAN: My name is Jason Tugman. I'm the VP of Cyber Risk Engineering at Axio. We are integrating a threat model ontology into our platform. That's actually an image of it there that we'll talk about shortly.

DR. VASSERMAN: I'm Eugene Vasserman. I'm faculty at Kansas State University. I'm also a senior staff fellow at the FDA.

MR. FITZGERALD: Brian Fitzgerald. I'm in the Office of Science and Engineering Labs involved in frontline review and some semi-support.

MR. COLEY: Steve Christey Coley with the MITRE Corporation. I help to support FDA primarily in the areas of vulnerability analysis, and I'm looking forward to this panel.

MR. CHAPMAN: Finally, my name is Joe Chapman, and I'm with The MITRE Corporation. I am a principal hardware security engineer and assist the FDA in cybersecurity as well.

DR. CARMODY: Thanks, gentlemen. So the folks sitting on this panel this way, this way, may have involved -- you may know them, they may have been involved in the guidance document and may have been involved with some of the things that come to our doorstep from your respective companies. Also Eugene over there on the end, I don't know why you sat over there. I'm not going to ask him to move. I think before, Jason, before I kick you off, you know, we have a manufacturer perspective as well, which is really important and then we have sort of an -- I'm going to call you an outsider, though, okay?

MR. TUGMAN: Sure.

DR. CARMODY: Okay, we have an outside perspective who's going to tee this up for us and get us started. So, Jason, why don't you get going.

MR. TUGMAN: Great, thank you. So a little background. My expertise is in cyber risk quantification. I work primarily with the energy sector, heavily in the energy sector, oil, gas, bulk electric, super majors, all through the pipeline security, etc., both in the U.S. and in Canada, as well as finance, helping them do very similar things. So what, for us, threat modeling is, is -- and I think that Seth mentioned it perfectly today and you said, for threat modeling, you really need to do and identify a function, which is, is the device connected?

Can you cause harm? But a threat model answers the third question to complete that cycle, which is if so, how? And that's what a threat model should do and that's what -- if you look at this ontology -- I'll skip the kind of forward bits that really just explains how to read an ontology. An ontology is a visual representation of a threat model. In fact, it is threat vectors. So on one side you have the threat coming in to your organization. On the other side you have a loss event or a loss scenario. I like to call them loss scenarios, not risks, because risk can be amorphous. This is a loss scenario causing harm in whatever way it can, so on that side.

So how those threats are going to enter your system, those are your vectors, and each prong identifies controls to prevent, detect, and respond to those aligned to the CSF, identify, detect, correct, respond and recover. So all of those relate. So this is actually -- this ontology was made by the open-source panel. This is actually part of the FAIR asset ontology; it's general asset ontology. Axio is developing, as Suzanne mentioned today, a stakeholder ontology for different sectors to be a little bit more specific.

So as you go through this ontology, you're going to obviate, you're going to remove the vectors that don't apply to you. But what this gets to is the last bullets here, is an ontology, a threat model, and an ontology is a way to do threat modeling. There are others, OCTAVE Allegro, etc., which is very similar to this. This is just the visual portion, in these last bullets, right?

Doing this activity gets to Section B of the cyber -- of the risk management documentation for premarket guidance, specifically B-1 and B-5. So what this enables, in closing, is that if a manufacturer can do this threat model, this ontology, the manufacturer can say here are the -- here is the loss scenario and it's many, here are the threats and it's many, here is our ontology breakdown of the threat vectors associated with those. These controls. So for each of these prongs, you're going to apply one or more controls, firewalls, segmentation, trust, no trust, etc., and a manufacturer can say these are our responsibilities. HDO, handshake, these are your responsibilities for the risks that we've identified. So this ontology or threat modeling, in general, facilitates the manufacturer device handshake to the HDO. That is the purpose, to me, what threat modeling is, it's identifying the controls and the threats.

DR. CARMODY: Thanks, Jason. Responses to that? Yes.

MR. CHAPMAN: I'll lead off. So Joe Chapman again. So yeah, I think this is a very -- I think this is a great contribution to the community and I think it helps kind of guide thinking when it comes to systems-level threats and analysis.

In preparation for this, I wrote down my own definition of a minimal viable product, which is one of the questions to the panel, and so I just wanted to kind of talk about that in the context of this contribution. So to me, when I'm doing threat modeling and think about the purpose, what's the purpose and what's sort of the outcome of this, I would say it's a well-informed and tangible understanding of adversarial-based risks to a particular system.

And so well informed, what that means, that means that you both understand the system that you're trying to model in the first place, as well as you understand some risks that are adversarial based, in general, so you've got experts involved, you've got proper levels of documentation and understanding on the team that's actually doing the threat modeling. And then the tangible piece means you've got some documented results that comes out of this.

So if it's in the form of an ontology, if it's in the form of like an attack tree analysis, if it's in the form of some other documentation, it's a TAM result now that can be shared for posterity, you know, communicated with other teams, communicated with peers for review, and communicated with the FDA. So that is kind of -- I think this is an excellent potential way to get there. Of course, there are many.

MR. FITZGERALD: I'd like to leverage off that. Threat modeling can be often seen as a blue sky and a little bit abstract. For the medical device manufacturers, they have to make a real threat model and it has to do with their intended use statement and it has to do with the context of use, their basic safety and their essential performance, which is likely to be prejudiced by intrusions. So it can be brought down to a manageable level if you think about it in the right way. There are techniques to arrive at the constellation of threats, other techniques that should be applied to limit them down so that it's a tractable engineering problem.

MR. SUÁREZ: And just to piggyback Brian's comments and give you something tangible, perhaps, to take away with your respective organizations. It may be even a question for the audience. You know, how many of you, by show of hands, have threat modeling integrated into your quality management systems today?

(Show of hands.)

MR. SUÁREZ: How many of you have a staff of over 10 product security/cybersecurity professionals within your manufacturing organization?

(Show of hands.)

MR. SUÁREZ: Okay, not many hands. But how many of you are manufacturers in general? And it's okay if you don't have lots.

(Show of hands.)

MR. SUÁREZ: All right. So you're probably going to need to go back to your boss, to your leadership team and tell them hey, look, I just went to the FDA workshop and I learned a little bit about threat modeling and it wasn't just a bunch of acronyms, right? But, you know, here's what we can do today with no resources whatsoever. You know, I oftentimes like to think of threat modeling as a brainstorming activity and you can take structured and unstructured approaches to threat modeling. You know, grab a whiteboard and start thinking about four simple questions, which is

1. What are we building?
2. What can go wrong?
3. What are we going to do about that? And
4. Did we do a good enough job?

And by the way, that is all documented in OWASP. OWASP has a great website with a threat modeling cheat sheet, by the way, that talks about that.

So I just said this is a whiteboarding activity and you're probably thinking great, but like I've got to put this into my risk management file and, you know, I need a template, right, and the template probably needs to be validated, right? And I think that's important to do, but don't lose grasp of having this be a whiteboard activity with as many stakeholders in your R&D organization and get them thinking about what bad things they can do to your product and again documenting that in a risk management file.

And by the way, if you don't have hackers, again, there's an enumeration of threats that you can use, something called CAPEC, C-A-P-E-C, as well as DREAD. Yeah, STRIDE, I'm sorry. I wouldn't recommend DREAD, but STRIDE, S-T-R-I-D-E. And those are really structured approaches rather than having like a free-for-all brainstorming session.

DR. CARMODY: So it sounds like we need a whiteboard, probably some markers.

MR. SUÁREZ: Some whiskey.

(Laughter.)

DR. CARMODY: Hey, I'll leave it to you.

Steve.

MR. COLEY: Yeah, what Rob said dovetails nicely into what I was wanting to talk about. Joe had mentioned the notion of threat modeling being a mechanism for sort of relying on being well informed about what your threats are and that's an area where things can become a little bit problematic. If you're coming at analysis of a medical device from a more or less pure safety perspective, you're not going to be well informed in terms of things that hackers can do maliciously. They will go and they will violate the laws of physics; they will go and take any of your estimations of likelihood of an event occurring and change that likelihood to 1.

And to be a little bit more well informed in some of these areas, besides acquiring your own expertise, as Rob had mentioned, there are a number of different freely available resources that are out there that can get you started, as well as proprietary solutions and consulted oriented ones. But it's a very cool idea, it's a very good approach, I think, to being a little bit more systematic and intentional in terms of understanding your medical device security risk.

DR. CARMODY: Thanks, Steve. I know that would mean you have postmarket issues. The first thing that we have started to do is construct our own threat model, which helps us.

Eugene.

DR. VASSERMAN: To follow up on what Rob said and then what Steve said, for those of you who do not currently do threat modeling, do you have a formal or informal definition of who your adversary is? Raise your hands if you do.

(Show of hands.)

DR. VASSERMAN: I saw two and a half.

(Laughter.)

DR. VASSERMAN: And one of those, I know for a fact, does it. So that's a problem. An attacker is probably the most important entity in this diagram. It is in the slide as threat. But keep in mind, this is an intelligent attacker. What Steve said, violate the laws of physics, they won't quite violate the laws of physics. They will appear to. Lightning can't strike literally at the same time in two distinct places, but an attacker can because they don't have to generate lightning, they just have to generate enough current to break down whatever components are redundant in those two places. So think like an adversary and an adversary controls the environment, so any assumptions you make about the environment you must state and you must enforce whether through -- labeling is not my favorite choice, but whether through labeling or through runtime checks, automatically or manually, what have you, that's extremely important.

The other thing I want to say is STRIDE was mentioned, but I want to emphasize the system-level thinking that needs to go into this. One thing that was left out, Jason, I think you left this out, and that is -- well, someone left this out -- no, Rob, you left it out. Bad Rob.

(Laughter.)

DR. VASSERMAN: What are your dependencies? On what are you depending to make your devices run? So for example, I make the monitoring station and I make the monitor. I don't make the network, so it's not my problem, right? Well, if the device is mission critical and it only displays alerts on the station and someone cuts your network cable because they don't like you or they don't like the patient, they don't like the facility, what have you, that is your problem. Your device, your system of devices, must have an active network. There are various ways to work with that, of course. The simplest is the alarm.

So you're not screwed, but that's one of the more difficult problems is networking out of service, but I'm using that to illustrate a bigger issue. Not only what have you built, but what are you assuming is going to underlie what you've built? What is the use environment? What are the other services provided? If you assume continuous access to the Internet, why and what happens if that's interrupted?

DR. CARMODY: Thanks, Eugene.

Jason.

MR. TUGMAN: So, you know, one of the goals of this panel was to identify what organizations can do Thursday when they leave this. So one of the things that I want to get to was go back to, kind of, Steve's comment about what is the MVP, what is the minimum viable product, right? And so to quote a 1945 mathematician named George Pólya, who's kind of the founding father of quantification, he said if there is a problem too large to solve, there is a smaller problem you can solve. Find it, right?

So the take-home is all of these, you know, things that we talk about are very important in the sense of the unknown unknowns, the Rumsfeldian unknown unknowns, there are so many of them. But there are so many, that's an unsolvable problem, right, but there are things that we do know. So the things that we can do is we can apply the basic questions that we've all been talking about kind of in the same way. What are our loss scenarios? What can happen to this thing, and it's many things, right? How can it happen? Eugene's point is so important. How can it happen? And that is to the device, from the threat vectors to the device or threat vectors or loss scenarios to the dependent assets. So a BIA, a business impact analysis, would include that, you know, the dependencies.

But to do on Thursday, go home, the whiteboard session, that is perfect. That's, in fact, how we do it as an organization. We get all the stakeholders in the room, whiteboard out all morning. We whiteboard out all morning and we pick a couple and then we choose to quantify those and that's a starting point that teaches them the methodology. You can do all of this on a whiteboard because it gets you to the true nature of quantification, which is the true nature of quantification is the reduction of uncertainty. That's really what we're going for, the reduction of uncertainty. So whiteboard what can happen, you know, what do I care about, what can happen to it, and then how can it happen. And then that gets you your threat vectors. That starts to define your controls. That starts to define your control handshake. That is threat modeling Step 1. Maturity after that, they're all the things that Eugene and others and Rob are talking about but Step 1, the MVP, that's it. And I don't know, I can't read that.

DR. CARMODY: Thanks, Jason.

DR. VASSERMAN: There's a set of cards called Elevation of Privilege that I'm going to plug, which are a great set of tools to start with if you're just whiteboarding this, and I encourage you to use them because they're fun.

DR. CARMODY: He'll provide a link later. Thank you, Jason. I can't deny you twice, you've been at the -- so we're going to have to go take your question, sir.

DR. ROY: Sabyasachi Roy, I'm the Director of Regulatory Affairs, a small medical device company called BrainScope. One of the challenges that we face here is that, as you all said, cybersecurity is a shared responsibility. The device manufacturer has to make some assumptions about what the healthcare provider is able to or not able to do. One of the challenges we've seen is we have customers who are large EDs, hospital systems that have more manpower than our entire company put together.

You talked about home healthcare. One thing that I haven't seen talked about here is the emerging urgent care market. We have urgent care to our customers who are somewhere in between, that you talk to cybersecurity and they all look at you and they'll probably blink and move on. So that's one thing, those are the assumptions of the shared responsibility and I communicate to them okay, here's your handshake, here's your responsibility. Are they actually able to do this, that's one thing.

The second thing is software as medical devices. Those things are you can build a software, perhaps, bill of materials, but they go on any hardware, are those hardware capable? How is the guidance going to deal with those two? I think those are very important because they're going to be sort of the face of the future and they are real for us, as a small medical device company, but I think that's going to become more pervasive. Those are the two points I had.

DR. CARMODY: Thank you. Responses?

MR. TUGMAN: So my thought on that is, you know, I work with pipeline companies and their margins are just absolutely tiny and they have some of the same questions and what I talk to them about is -- and you talked about software, right? So one thing about trust, you have to look at zones in trust, right? Trust on trust. But trust is also zone based, so if something is un-trusted but is in a network that is trusted or you have the communication flow between or a low-risk device, as we talked about in the previous panel, and a high-risk device, if those are in the same zone, that low-risk device has to be elevated to a high-risk status, right? It's trust, that's zoning.

So when you look at how to do that handshake, the thing that we talk about, Step 1 for the really small folks that have 0.5 FTEs of cybersecurity, right, but they deliver gas to grandma, to 500,000 households, right? That is a real case in many, many instances. We begin with the loss, what can go bad. We kind of skip the best practices and we align the loss and then say here are the best practices. It's actually the middle bullets there. We align the practices to CSF, C2M2, whatever practices, to those loss scenarios specifically. So it's not apply best practices, it's apply best practices that relate to the loss and it focuses them and it allows you to have a more directed conversation about control, not just are you doing access control. That means nothing, right? So that's how you begin those conversations of trust.

DR. CARMODY: Rob.

MR. SUÁREZ: I would just add to that -- I think that was a great summary, and I would just add to that the mindset that you have going into risk assessment or threat modeling and even for our penetration testers, and I think that's a given for our pen testers, is to have a worst-case mentality, that if it's possible it will happen, that worst-case scenario. And I think the reason why that mindset is helpful during threat modeling is because there is a likelihood wormhole where you can spend an exhaustive amount of time trying to think about can this happen and maybe not, you know.

The other thing I would invite you to do is to reach out to security researchers, to reach out to your customers and invite them to your risk assessments, to your threat modeling activities, and I think having more people involved in those types of activities provides more fruitful discussion. I always think of cybersecurity risk and risk in general, like a room with many windows and many different corners and pockets and every time you look through a different window you can map out a different part of that room.

DR. CARMODY: Thank you, gentlemen.

MR. FITZGERALD: Just to add one more thing, if I could.

DR. VASSERMAN: Could I --

DR. CARMODY: Eugene, please.

DR. VASSERMAN: Sorry, I was waiting patiently, which is very uncharacteristic of me.

(Laughter.)

DR. VASSERMAN: So I just want to touch on two things, which is if it can happen, it will and also Rumsfeldian unknown unknowns. I'm the token academic here, I believe, so I'm going to say the F word: formal methods.

UNIDENTIFIED SPEAKER: That's a beautiful process.

DR. VASSERMAN: Eventually it would be good to build up to formal methods. You don't know what the unknown unknowns are, but you can list exactly what can happen in your system and you can prove, mathematically, that nothing else can happen given that certain assumptions about the environment, your hardware and information theory hold. So while we're not there yet, I believe we can get there and with some effort, we may even be able to get there within the next 5 to 10 years, but of course being faculty I don't know what I'm doing, but I think I know what everyone else should be doing. I'm still getting fired for this.

(Laughter.)

DR. VASSERMAN: But anyway, please, if you're not working towards formal methods, baby steps, just as much as you can show formally. And this ontology is actually a great step in the right direction. It's a very informal approach to beginning a formal methods process. Please, please do that. It will pay off tremendously.

DR. CARMODY: Thanks, Eugene.

Brian.

MR. FITZGERALD: So the idea of threat models being static should be thrown away immediately. Medical device manufacturers in the real world will want to sell their devices to as many people as they can and that means that new types of threats have to be introduced into that overall threat model which means the design must iterate as new opportunities, business opportunities, come along.

It is therefore reasonable to ask the question if you -- and our last questioner was asking about urgent care facilities. It's reasonable to ask have you included urgent care facility based threats in your overall threat model, are they any different from your threat model as it exists? Would that proscribe you from wanting to sell into that community, if you didn't know what those threats were? Do you have processes to control the way that your threat model drives your business process?

DR. VASSERMAN: Brian, surely those are use cases rather than threat model.

MR. FITZGERALD: But it could be that they are very specific and unique threats in the point of use and the context of use.

MR. TUGMAN: So just to pull on that, right, just to pull on that thread and I apologize. So when you look at -- that's part of the handshake. So what Brian's mentioning there, it's part of the handshake. You, as the device manufacturer, once you go through this threat modeling in whatever form it takes, you know, the attack trees, which is like an OCTAVE kind of way to do it; other ways, OWASP, to do it, it's also an attack tree, I believe. So once you do all of that, you, the manufacturer, will know what controls you're responsible for.

But also we talked about dependencies. Controls are also dependent, right? Not just the hardware or the environment, controls are dependent. In fact, this ontology stacks because defense in depth stacks, right? So that can go five layers out because defense is in depth. So as you, the manufacturer, are going through this process, you know what controls are critical and if an environment, if a use case is at a nascent state -- that's a polite way of saying can't -- you know, is it a nascent state in a cybersecurity capability, then you have to therefore recognize -- and in the risk sense, right, so you can accept the risk, you can mitigate the risk, you can transfer the risk, right, you can avoid the risk. So you can either avoid the risk or you can accept the risk, those are your two choices because you've done the homework and you can make some heuristic analysis of the environment use cases and then you're then informed about the risk that you are accepting into your environment, which could be legal acceptance as well, right? Financial acceptance.

DR. CARMODY: Rob, Joe, question?

MR. CHAPMAN: Yeah. So I just want to dovetail real quick on, Jason, some of the comments you're making and the question about software medical devices, I think that is a really important question and I think, as a system designer for that software device, I think you need to put requirements on the hardware and have certain assurances that are provided by the hardware vendor and that is absolutely part of, and within scope of, the threat modeling exercise that you do. What if those assumptions are violated in certain ways, how does that actually impact my device? And so, you know, I think -- you know, do you start with the requirements and then think about what happens if they're violated or do you start with thinking about the threats and then that drives your requirements? I think it's very much this kind of circular iterative process. I tend to think with the threats because in my mind, I like to think more on the analysis side of things than the design side of things, but the two are intertwined.

The other quick comment I want to make is we've talked about the adversary and thinking about your threat from the adversary's perspective. I want to say that that is a space where we get into thinking about human behavior and modeling human behavior and there's societal and other elements to that that I think are very much so on the unknown category, at least.

And so I think focusing your efforts and focusing your thinking with a few arch types of adversary, right, so a few kind of threat actor templates, let's say, with perhaps distinct objectives, maybe some have financial objectives, maybe some have damage as their objective. Think about them in the abstract, and then focus your efforts on your system, on what you can know and what you can control, I think is an important aspect. I wouldn't want a lot of effort wasted, let's say, on trying to think about specific threat actors, you know, in the context of this exercise.

DR. CARMODY: Can I ask you a question? That's not the way it works, right, you ask me -- okay, all right. I had a question, but okay, go ahead, sir.

MR. WILSON: Charles Wilson, senior architect at Draeger Medical. When we --

DR. CARMODY: Closer to the microphone.

MR. WILSON: Closer. Charles Wilson, senior architect at Draeger Medical Systems. When you look at threat modeling and you look at the -- I guess the state of various manufacturers, given that there hasn't been previous guidance, it can probably be assumed that threat modeling is going to be new to a lot of these people. And my past experience with threat modeling -- I'll take this from a software perspective -- is that a particular threat model will include roughly one threat scenario per thousand lines of code and that can present a larger amount of, if you will, surface area to look at and the amount of time that it takes create is non-trivial.

When this guidance or when this guidance comes out, it would probably be very helpful, especially for the companies that have no background in threat modeling, to give them a sense of what the scope is, what the amount of commitment that's going to be required is. Also, the distinction between creating a new device and creating a threat model for it versus submitting a device update which may add a feature to it and what impact that's going to have on them. Are they going to be required to generate an entire threat model for something that's been in the market for years for the addition of a single feature, or will it be acceptable to have a feature-specific threat model created which is scoped much smaller than that?

Additionally, you have the issue that goes along with this, that was brought up earlier, of one of the things that's attempted to be done is de-risking things by pushing them out, either whether that says we have a device here, a device here, and then we have the nebulous miraculous network and to a large extent people will say well, that de-risks that for us. However, that doesn't mean that there isn't a threat present there and it doesn't exclude the necessity for a threat model for that. Well, that presumes that that company is cognizant of what a threat model is in the first place and second, can produce one. And even if they can produce one, how long will that take? The impact to organizations which today are producing things and which have produced things for years is going to be fairly significant, especially when you add the thing that everyone who does threat modeling understands, which is that this is a not one-and-done --

DR. CARMODY: Right. Sure.

MR. WILSON: -- that you will have to do this every cycle. And if you're smart, you're combining this within your attack surface analysis, which, for embedded devices, is more or less problematic today.

DR. VASSERMAN: And if you come up with all of this information, how is it problematic for embedded devices?

MR. WILSON: I consider it problematic with respect to desktop operating systems where you can -- where you can have things which are profiling the system in the state that it's in, in terms of a pristine state and then an after state and being able to compare the two of those, being able to look at all of the port activity that's going on, all of the operating system calls which are made, all of the various interactions that are going on in whatever represents a file system. With a normal operating system, that's a fairly trivial operation.

MR. TUGMAN: I can name three technologies that do that exact thing with embedded devices. There's also an architecture model that we use in the energy sector called the Purdue model and it takes tiers. And so you would be talking about kind of Tier 1 or Tier 2 device, and in between each of those tiers are trust zones, you would have DMZs, you'd have multiple DMZs per to your -- that's very advanced, right?

And so when we're talking about what we can do tomorrow, what we can do Thursday versus what you're talking about, which is very advanced, very deep and very important, but when we talk about on the maturity scale, right, the people in the room, there were seven people that raised their hand that they even have a team to do threat modeling. So the MVP, going back to Brian's point, let's -- you know, focusing there on what can we literally do tomorrow with our teams when we get back to the office Thursday afternoon, right?

MR. WILSON: Absolutely. But I would appreciate if that could be folded into the guidance and not only give them guidance of this is what we expect, but also these are steps that will help you get there.

DR. CARMODY: So just to tease out a couple of the things you said are very important. They're complex issues. Please submit comments to the docket. I see some of your fellow Draeger folks here, so please do submit comments to the docket. You're welcome to come in. We will supply the markers and we'll have a whiteboarding session. We've done this. Like the question is so big and it's new to some folks, let's have that conversation. So happy to do that. In terms of new/old, we have to make a call, okay. Okay, and it's going to -- in some cases it's going to be painful. Let's have a conversation about when we pull the band-aid off in terms of when an old device gets an updated threat model and there are some things that need to change with respect to that threat model. So I'll pause there and kick it over to Joe.

Steve, go.

MR. COLEY: To continue on this particular topic, you touched on more or less how much detail would need to be provided, how much detail is important, and alluding to the expense, effectively, and the amount of time and personnel that may be needed to provide that. That's one thing that I think I would say a number of us here on the panel are sort of dealing with in terms of ongoing activities is how much information, how much details do we ask for? How much is it reasonable to expect?

You don't want to demand too much of a manufacturer. On the other hand, we don't want to demand too much of a manufacturer and then receive more than we can process, either. And it's an ongoing difficulty, but I think in some of these recent -- recent times that we've been fortunate to have some ongoing dialogues with some manufacturers to help try and sort of tease some of that out. I would suggest, and I might be wrong in this, but I would suggest that adoption of threat modeling tools and associated, kind of, capabilities may be able to help with some of those sorts of requirements by effectively producing some of the documentation automatically as a result of conducting those kinds of exercises. That might not be correct, but that's a suspicion I have.

MR. CHAPMAN: With regards to the challenges for embedded devices, so I take that to be synonymous with resource constrained, right? So the example that you brought up was in an operating system, you know, on a desktop class machine, I have a lot of processing and resource overhead in order to do kind of continuous integrity monitoring, right, runtime integrity monitoring. That might not be possible on a resource constrained real-time kind of system. What I'd say to that is, is that the threat modeling exercise is still valuable. I'd also say that it's necessary. Basically, it's a necessary exercise to go through to think about this at the system level, not just at the device level.

So at the system level, what's the -- so if I can't protect this thing adequately or if I can't protect this against all of the threats that I've identified in my model, what's the consequence of this thing being fully compromised, right? Is it isolated to a single patient that has already been, you know, within physical proximity of an attacker? Is this that I've lost some core secret now that the rest of my devices are now vulnerable, right? So what's the consequence of that compromise? And I think that's very much all the way at the left side of the chart here, right, the loss event, the risk that's associated. I think taking that perspective and making arguments and reasoning around that, I think, is pretty valuable in this exercise overall. So it's not necessarily let's make the objective impenetrable embedded systems; it's let's think about the risk at the system level and how this fits in.

DR. CARMODY: Eugene.

DR. VASSERMAN: Just a brief note. And I speak for no one other than myself. I want to point out that the threat and risk to a patient is real, whether or not it's coming from a security vulnerability or from a manufacturing defect or from material degradation.

So while I recognize that it may be very difficult to start thinking along these lines, especially if you've never done it before, it is critically important for the patient. They are just as hurt whether they suffer as a result of WannaCry or from toxins leaching out of a plastic.

MR. SUÁREZ: And actually a really nice transition to mention when thinking about threat modeling, I also -- I think I'm probably talking to a room full of mostly medical device manufacturers and regulatory and quality professionals, so I'd imagine there's an appreciation for process and formality behind this.

If you go back to your teams, you know, have a discussion around perhaps how threat modeling can -- has a lot of parallels to hazard analysis, which we oftentimes do today as medical device manufacturers, but from a clinical and safety perspective and perhaps from, you know, incidental or, you know, non-adversarial types of events, right, but also that the end product should not be a threat model for your customers and for the FDA. Don't get obsessed with the process. You know, focus more on the results, which is that threat modeling can produce a more comprehensive set of risks, a description of the risks, the threat actors involved, explaining those risks to, yeah, FDA during a market submission but also to your customers when they're buying this product and you need an explanation as to why I need to pair my device within a certain proximity. You know, the end result. Again, it's just a higher quality description of risk to your patients.

DR. CARMODY: Yes, 30 seconds and then we'll do final thoughts.

Jason.

MR. TUGMAN: So I want to put a button on everything. One of the things -- so you mentioned hazard analysis. One of the things we do in the energy sector is we -- what I do with my clients is I look for what's called inflection points, right, what do they do well? So the energy sector does physical security extremely well, right? They've done it really well for decades. One thing that the medical world does really well is hazard analysis. So as you go home, look at the processes, we're talking about repeatable processes, look for the repeatable processes that can be inflection points into this very robust -- it doesn't have to be -- you know, the end state is extremely robust, but the beginning state, we say, is -- you know, it can be ad hoc. Something is better than nothing. Look for the inflection points in your organization, and I do agree that hazard analysis is probably one of those inflection points.

DR. VASSERMAN: And, ideally, hazard analysis would catch security problems which terminate in a patient -- patient risk.

DR. CARMODY: All right, folks, we're getting close to the end, so final thoughts? I'll wait as long as possible until somebody rings in. I'm looking over here, I'm looking over here.

UNIDENTIFIED SPEAKER: We've talked enough.

DR. CARMODY: Good end. Over here.

MR. COLEY: We haven't touched much on this during the panel today, but one thing that's important, whether it's threat modeling or other techniques for conducting cybersecurity risk assessments, is really look at the device as one part of a larger system of interconnected components, some of which you, as a manufacturer, may also own and control and part of that analysis really does need to consider what happens if you lose control of one of those components. Say it may be compromised by a separately independent vulnerability. It might not touch the device directly, but it might be able to be exploited in a way that does impact how the device operates and therefore has patient safety considerations. So I leave with a small suggestion here: trust no one, including yourself. Or us.

UNIDENTIFIED SPEAKER: I agree.

DR. CARMODY: Highly unusual for a panel to say that, but you know you've got good security folks when they say that. A question to the back.

MR. NITCHY: Yes, Carl Nitchy (ph.), Canon Medical Systems. This appeared to be quite a bit of good useful information for the HDOs themselves, so they now have well-informed information about the threat modeling as we've done it. I've seen HDOs ask us if we do threat modeling. From the HDO standpoint, is this something that you're going to be requesting or asking manufacturers to provide? Similar to the MDS2.

DR. CARMODY: You're saying have manufacturers provide a threat model to hospitals or their customers?

MR. NITCHY: Correct, because there's obviously quite a bit of sensitive information there for a company. It shows the vulnerabilities and risks inherent to the system.

DR. CARMODY: Yeah.

MR. NITCHY: However, it's really important for the HDOs to know those risks.

DR. CARMODY: Sure. So I think the thing that we have in our premarket guidance draft is emphasizing transparency. We haven't specifically said a threat model, but we have laid out a number of labeling items that we think get to the transparency issue. I think the conversation could have many ways and that's probably an open question, how do you effectively communicate your security posture to your customers?

MR. NITCHY: Right.

DR. CARMODY: Final thoughts.

MR. TUGMAN: Sorry, Rob.

DR. CARMODY: All right.

MR. TUGMAN: So one thing --

DR. CARMODY: Five seconds.

MR. TUGMAN: So if you look at this, it says B-1 and B-5. It actually should be B-2. It gets to your question exactly. So in the premarket guidance, it actually covers what a threat model could look like. Joe had mentioned the importance of doing threat models. At an AdvaMed conference we had a couple months ago, one thing that came out of that is standardization. So if we look Item 2 and Item 5, which is on page 23 of your premarket guidance, it covers a lot of what we just talked about today and it really breaks it down into a simplistic form.

Lastly, Rob had mentioned the likelihood wormhole. I'm just going to give a prop, the shout-out to someone who battles with this topic all the time, being -- especially being quantification. I don't care about likelihood or probability. That's weird to say. I start with impact and if you're going to do likelihood, and I hate that I'm going to say this, use a high, medium and low. Don't go probability, you're going to run around and it's going to get you into the likelihood wormhole. It was a great comment. And yeah, Item 2 specifically says exploitability as described, likelihood instead of probability. If numerical probability is provided, we recommend providing additional information to explain how the probability was calculated, right? That is so critical because most times it's fuzzy math, it's fake math. So it's awesome it's in there. Sorry.

UNIDENTIFIED SPEAKER: I would like to agree.

DR. CARMODY: What a beautiful segue to the next panel. Folks, lots of good conversation happening up here, it's not the end. Please thank our panelists for providing awesome information. Thank you.

(Applause.)

DR. D'AMICO: All right, we'll move on to our last panel before we break for lunch. So if you're a panelist or the moderator for Risk Assessment Approaches and Labeling, you are welcome to the front of the room.

DR. CARMODY: Ladies and gentlemen, I have some unfortunate news. They've retained me for the following panel; I'm very sorry.

(Pause.)

MR. ROTHSTEIN: Okay. Well, welcome to Session III, Risk Assessment Approaches and Labeling. Seth is still up here. We have transferred from one bald head to another, so for those of you in the back, it might not look that different. My name is Zach Rothstein. I'm a vice president at AdvaMed, the medical device trade association, in our technology and regulatory affairs department. And I just want to start off by saying thank you to everybody at FDA, Suzanne, Seth, Aftin, Reid. They have done a tremendous job putting this event together, and in particular, the last of a month which was during a government shutdown. So congratulations on putting together today's event.

(Applause.)

MR. ROTHSTEIN: And the turnout clearly, you know, shows how well the agenda came together. Let's just begin by going down, starting with Dana-Megan, with bios and introductions.

MS. ROSSI: Good morning, everyone. I'm Dana-Megan Rossi, and I'm the Associate Director for Product Security Operations at BD.

MR. HOYME: I'm Ken Hoyme. I am the Director of Product Security at Boston Scientific, and I will note at this point two out of three panels have been led by a person named Zach.

MR. FISCHER: Hello, my name's Christoph Fischer. I'm a lead system engineer for an insulin pump system at Roche Diabetes Care and co-chair of the IEEE PHD Cybersecurity Working Group.

DR. CARMODY: Seth Carmody, Cybersecurity Program Manager at CDRH.

MR. BITZA: I'm not going to talk that fast. My name is Chris Bitza. I am the U.S. product security leader for bioMérieux.

MR. MORGAN: Colin Morgan, Director of Product Security at Johnson & Johnson.

MR. ROTHSTEIN: Great. So the purpose of this panel is to get into how companies assess risk with a medical device throughout the product life cycle and also labeling challenges as well. So let's just start. Maybe, Seth, I could ask you, could you just frame this for us, between what is a Tier 1 and a Tier 2 device under the draft premarket guidance?

DR. CARMODY: Yeah, absolutely. So just to repeat some of the remarks I had in my initial talk. So the tiering system really reflects a couple of things. It's a risk-based approach and it's also intended to be least burdensome to identify for you folks, through a threat model, the things that we've seen and then to suggest what are the controls for those risks. By default you're Tier 1 and then you have to sort of argue your way out of each line item as to why you might be Tier 2. Happy to have conversations around that as well, but that's probably the most concise version.

MR. ROTHSTEIN: Okay. And maybe, then, before we get into the risk assessment piece of the panel, do any of the panelists have any thoughts about the two-tiered system that they'd like to mention?

MS. ROSSI: I think Ken and I were both jockeying for the microphone. Yes, I would ask us and propose that perhaps we take a vigorous approach to risk assessment and a life cycle approach to product security regardless of whether something may fall into Tier 1 versus Tier 2. To that, we're always applying best practices across the board for all of our products.

MR. HOYME: I think I'd reach the same conclusion with a slightly different observation as in looking through the controls that are mandated as a Tier 1 device. I will provide in our feedback examples of situations where some of those controls, for risk reasons, would not be appropriate which, really, is the same argument you would do for a

Tier 2 device. So I really think, from a risk management approach, what that list that is in Section 5 should serve is as a guideline of. All of these should be thought about and, by default, we'd expect them unless you can provide a risk reason or use environment reason why a particular control might not be appropriate. So I think there will be lively feedback on the tiering system.

MR. BITZA: Yeah, and building off of what Ken just said, when you look at what declares the Tier 1, there's an "and" statement, right, they're connected and multi-patient harm and once you have the first point the "and" might be better replaced with what follows "is," that we know when we -- you know, building off the device panel with the threat modeling, so many vulnerabilities might default once it's connected or multi-patient impact, in which case it's a single-tier system.

MR. MORGAN: I would say that FDA has successfully, two times in a row now, added a very controversial topic that is very hotly debated. To go back to the postmarket, the 30/60-day remediation window was a hot topic, just like everywhere I go and everyone I talk to, Tier 1 and Tier 2 is the only thing anyone wants to talk about.

What I would say is that, you know, it's the concept of taking a risk-based approach, I think, is key and very paramount, but we also have to be careful to not remove some of the ability of a manufacturer to follow their existing risk-based methods to fall into this different system. Like I mentioned on the previous panel, you know, devices can range, there's a broad spectrum of things that may or may not be a device, a mobile app, a software platform, to these invasive technologies and we have to make sure that we're not putting too much burden on the lower-risk devices but at the same time making sure that we're doing right by the higher-risk ones.

So, you know, a recommendation from our side, and we'll have some of this in our response, too, is really just around figuring out how we can appropriately come up with that risk-based approach and to factor in these different scenarios.

MR. FISCHER: And I would like to raise a question, a pertinent question, with regard to what is meant with multi-patient harm, but -- thinking about an insulin pump route of connected or maybe a vulnerability, but you have to go in the range of all of these patients in all the -- compared to a network-attached infusion pump. So what is this direct meaning here? To multi.

DR. CARMODY: Yeah, so I think you sort of teased it out there. So direct harm, I think we put insulin pump or infusion pumps in the guidance as an example of a Tier 1 device. The multi-patient, if they present a multi-patient scenario where you -- a single vulnerability could affect all of the fleet of the devices. So the situation that you laid out where an adversary may have to go and modify each device specifically, you know, that's not what we were intending in terms of hitting, in terms of multi-patient harm, that an adversary physically or, you know, in some way has to know information about each of those devices to harm individuals one at a time.

MR. MORGAN: What about a device that is your jump point into an entire hospital? That could lead to more patient harm.

DR. CARMODY: Yeah, a good question. I don't know if we explored -- I don't want to use the word tier but yeah, as you get primary, secondary, and tertiary away from the device itself, what types of harm can that item or that device inflict? We did, in the original premarket guidance, talk about provide access points to a network. So if you, you know, have some weaknesses and you allow going into a hospital network, what -- you know, you have to think about those. In terms of the Tier 1/Tier 2 system, I think we've been silent on it. Do you think that we should say something about that? Is there a concern there?

MR. MORGAN: Well, I think just the way that it's worded today, it's a little bit open for interpretation, and how you just described it may be a better language to solidify that.

MR. HOYME: So I think we've had a broad discussion going on a while in this industry about how we ensure devices are good citizens on a network. Can they be used as pivot points? The FDA, from a regulatory authority, is looking at a specific behavior of a single device at a time. Hospitals are certainly thinking about it from a risk assessment, how would they assure that their networks stay up in the presence of threats that are on the network. So certainly we see contract language that speaks to that, we certainly saw with WannaCry and some of that, and I think the reference is, is availability as a potential patient risk, that if all you do is coordinate some kind of denial service to the attack that makes the devices unavailable, that could be a source of patient harm.

So I think there's still a struggle of how do we make sure that we are addressing the challenges of not making a device the weak link into a system and, in parallel with that, does that make it a weak link in terms of that device performing its essential performance for what the device itself was intended to do? So it's still a challenge in terms of how we document, even though I think, as an industry, our intents and understanding is that we shouldn't be the weak point that causes the network to be attacked.

MR. ROTHSTEIN: So I think it's safe to say, based on that conversation, Seth, that AdvaMed's comments will probably reflect something along these lines, since I think --

DR. CARMODY: Yeah, we really love -- just realizing the configuration of this panel is very overwhelming, so --

(Laughter.)

MR. ROTHSTEIN: So let's then move into the risk assessment piece. So open question to anybody on the panel. How are your organizations thinking about integrating the risk assessment output into new device designs? Anybody want to start with that?

MR. MORGAN: Yeah, I can jump on that one. So I always joke with our team that, you know, we don't do risk assessments, we do security engineering and risk assessment is just a part of that. Yeah, we heard about threat modeling. To me, threat modeling is just a way to analyze your architecture, your design, understand where your weaknesses are, identify what controls you need to put in place. And that is inclusive of not only your solution in itself, but also the interfaces it may have downstream, whether that is connected to an instrumentation, a cloud environment, a database, a server, however that architecture might be set up. And so, you know, we look at it holistically of understanding those designs, what they are and what needs to get built into it and then based on the risk of the solution, the controls get adopted appropriately.

From the risk assessment piece, you know, traditional 14971 risk assessment piece, we -- you know, we live by the model of continuously improve based on our own internal learnings and by, you know, collaborating with our peers in the industry and learning what they're doing. But we look at it from that perspective of we build out what our profiles look like and have things like if you're a CVSS medium or above, you're required to be remediated prelaunch. If, for some reason, that solution does not remediate one of those, it gets pulled forward to the safety risk analysis and identify there's a hazardous condition and potentially a patient impact could be evaluated for what decisions get made. So that, to me, the risk assessment piece is kind of a bit of an output from just your overall good practices around security engineering and architecture.

MR. HOYME: So, yeah, we have similarity. Depending on the business, we have had an integration of security risk assessment into our quality system as a parallel to the 14971. So if you see what's articulated in TIR 57 of AAMI, of how the parallel analysis processes would crosslink that, make sure that safety and security are integrated, and we started doing that on some of our products starting in about 2003 when we saw, you know, the concern of multi-patient concerns being a potential safety risk. So, you know, the real lever for that in a medical device product company is to do that linkage to safety. You have to enhance that because you could get in an overly strict interpretation, confidentiality, privacy issues, which you are also obligated to provide to your customers, that don't fall under the traditional FDA regulation authority of that. But if you want to sell into the marketplace, your customers are obligated to meet those as well. So we've looked at ways in why we separate our security processes to make sure that we are robust to not limit to safety impacts but in addition to safety impacts, look at privacy/confidentiality issues as part of the design process.

MS. ROSSI: Yes, and I'll follow up on that. At BD, we have a process where risk assessment is an essential part of our product security framework and, in fact, we have -- through the spirit of transparency that Dr. Schwartz mentioned earlier in her remarks, we have actually published our policy procedure, and you can see our product security framework that includes risk assessment and all of the different activities that are an essential part of all of our product reviews.

And so I invite everyone to take a look at that, as well as I would propose everyone take a look at the Joint Security Plan that's been mentioned a couple of times here. It's a tremendous toolkit that it allows different stakeholders within the industry to take a look at what are the essential elements for risk assessment and how does it form an element as part of your life cycle strategy. So as Colin said, we heard about threat modeling earlier today and we're talking about risk assessment as components to your overall strategy.

MR. BITZA: So one thing I would add, you have the opportunity to not have security just be the policeman in the room. You have the opportunity to work with your business partners when you're talking about adding security to the design of the device and perhaps you can bring technologies or solutions to mind when you understand their use case that they might not have otherwise have been considering, in which case, again, you're no longer the policeman in the room, the bad guy who just tells people what not to do or how to add time to their project schedules, but actually bringing solutions to the table for them, yeah, up front, early stages of the design of the product.

MR. FISCHER: And I would like to add something else from a different point of view, from a standardization point of view, because if -- for all of this discussion and maybe you have the same feeling. People name the same thing but have different meanings on it, so there are a lot of varieties behind it. And I think there will be an issue also coming, not just for industry making submissions where they have to make up their mind how they're doing it and maybe identify how they're doing it. Also, I think there will be issues for the FDA by going into the context, what is the meaning of that, and this is why I'm thinking that also this may be the next step.

Also, we should think about how to standardize some of this topic, now speaking from a -- point of view because I think this definitely can really help everybody do this and the efforts as well, also, to come to a better result because if you're thinking about the first version of the guidance -- there are about six bullet points to address. With the new one, we have about 30.

So for many companies, submission just means a lot of paperwork and I think, during this transition phase, maybe some could lose to focus on these whiteboard questions, which you had in the session before, so what is really the meaning on what we are doing here by trusting the paperwork. So I think by standardization, maybe also applying tools which help us, we can again be focusing better on what is necessary to do.

MR. ROTHSTEIN: So a slightly different take, then, on this same question. In terms of assessing risk of a device, we just talked about it from a new device perspective. How do the companies think about it from existing devices, right? And this comes back to the issue of devices that need to go back to FDA for an updated or new 510(k) or some other type of pre-submission request. Does anything change when you have that device that's already on the market and you're reevaluating it with respect to that new submission to the Agency?

MR. HOYME: So I think the challenge that you have is there is a classic watchword of what Ken bolts security on. There are certain aspects of -- certain controls and aspects that you would roll out are fundamentally cored to the platform you develop on. And so I think, you know, as an industry, and I don't think this is at all unique to the healthcare medical device space, there are iterations where you do platform updates where you may add hardware, TPMs, things of that nature. There are iterations where they are software based.

And I think the challenge is you need to go through, with the new guidance, the kinds of risk assessments that are appropriate to the security environment you're selling into, but there can be constraints on what controls and solutions are available within a non-modified underlying platform as a pure software update versus at what point do you step forward and have to produce an updated platform that has other security controls and capabilities built in.

And I think that's all in the risk process that companies go through, that purchasers go through, that I assume the FDA has wrestled with, which is when is something too long in the tooth to integrate more security to and fundamentally insecure and platforms need to change or not? But, you know, clearly the indication with the guidance, what we end up doing when we roll things into our quality system is, you know, incremental updates that are going to be required have to go through that kind of process for the assessment.

MR. FISCHER: You asked from the point from submission and resubmission. I would go a step further. At least in my area, we have something called a service security maintenance plan, so threat models, first get a license on a living document. And so, over time, independent from adding new features or bringing it to a resubmission, we go through it again and again on a regular basis and based on that, we make, maybe, adjustments, which maybe resides in an update.

MS. ROSSI: Yeah, I'll add on to that. When we talk about risk assessment, it's a point in time, right, and what we're doing is we're assessing the risk at that point in time. What happens after a risk assessment? Remediation, right? So it's an ongoing process, it's a journey, it's not something that you do a risk assessment and then you file it away and call it a day. This is an ongoing process so that you're taking the lessons learned, building it back in and continuing to improve upon your product.

MR. MORGAN: Yeah, the point I was going to make, well, like Dana just said, is a key word is learn. I think we need to not only learn from our own industry, but from the tech industry, what has worked for them and what has failed. There's been a lot of advancements in software enabled solutions that are outside of healthcare and they figured out how to do it, they figured out how to make $50 devices that they can patch whenever they need to and they figured out how to evaluate those risks and manage them properly.

Now, our sector is a bit different because of the safety impact and the patient impact but, at the same time, we can, you know, kind of pull those theories and methodologies forward and build them into our processes. That's where, you know, we spoke about the software maintenance plan, that's something that's part of your premarket design and you need to start thinking in an agile mindset and figure out if these devices are going to be updated, they're going to have to be maintained, there's going to be incremental versions.

You know, it's not like your -- an app on your phone gets a major version every other month. You're getting incremental updates, new features, new bells and whistles and some of those are security patches, some of those are security updates, some of those are, you know, new access controls so you can do face authentication. And in the device world, we just need to figure out how do we incorporate some of this into what we're doing but do it within the constraints that we have.

And you've heard a lot about safety. You know, safety is a key element here. Partnering with the medical safety organization ensures that you're properly not only meeting regulatory expectations around, you know, your 30/60-day timelines and understanding if there's safety impacts, but it's also about making sure you're doing right by your patients because, at the end of the day, these devices are serving people and if you're doing your best, then, you know, that's what we all hope for.

MR. HOYME: I just want to make a quick point. We drifted. We started with a question of if you have a legacy device that does have these capabilities, what do you do? And we drifted to and if I'm developing a new platform that has to live in this environment, what things do you need to do to make sure you do it? So absolutely what Colin was saying is when you start with a new platform, what are the capabilities for updating and patching and how do you assure the authenticity of those patches and those type of stuff really should be considered. The challenge always is if I'm trying to software update a device that didn't have those underlying structures, when is it -- yeah.

MR. MORGAN: I was trying to not pull us back into legacy.

MR. HOYME: Yeah. Yeah, we'll leave that for the --

DR. CARMODY: So, relative to your question, Zach, one of the things that we noticed when going through some postmarket issues is if there was a lack of security, no security controls whatever in the postmarket, we couldn't -- we weren't even in a place in the discussion to talk about safety assessment, safety risk assessment. So the conversation shifted from is there a safety risk here to go fix -- go do security controls and then come back and we'll talk about safety risks. And that gets into threat modeling and it's a more complex conversation, but that's one of the outputs that we had in the postmarket sense.

MR. MORGAN: And for those who haven't gone through it, a safety risk assessment from a security standpoint is no different. Update your FMEAs, update your PIAs, your health hazard evaluations, and just incorporate the security and try to pull forward the exploitability versus impact and leveraging, you know, the current processes you have around the probability values and the severity of impact values and it just -- you can fit it right in there with the safety piece.

MR. ROTHSTEIN: The next question I'd like to move to with the panel is about the challenges and opportunities in taking a systems-view approach to risk analysis, right? So this would be when you look at components outside of the device, such as your network, an HDO network, mobile apps. So, again, what are those challenges and opportunities when you take that type of view of your risk analysis?

MR. BITZA: Well, it builds off of what the former panel was talking about related to trust zones and you should -- because we know the cliché is once you've met one -- or seen one hospital, you've seen one hospital. Most of those things outside of your device you need to consider as already suspect.

Another thing that the -- building off of what the panel was talking about before about know your adversaries, a lot of inclination about the APT, the nation state, highly effective malicious user, your end user is your biggest threat. You need to assume that your device is going to be in a suspect environment from the go.

MR. FISCHER: As a system engineer by heart, the system is the starting point and especially here, if you're familiar with, for example, a system and a modeling language, the system context here, you can see it's a similar thing, a starting point for your threat models with the trust zone because the boundaries of your systems -- I'll first point out a few there -- you have to look what's crossed over this area. And then if you dive and decompose the system, looking at what happens inside and what happens there and could someone in this environment have access to the connection in between. So, yes, this must be the starting point.

MR. HOYME: So I'll reinforce that, what Christoph said was -- I also came at this from system engineering and security, like safety, is an emergent property of a system and you really cannot answer the question about whether a device -- a component is safe or secure without understanding the context that it lives in. And so the variations in hospitals, the variations in a device being given to a patient that's out in the mobile world, ambulatory with it, all have to be considerations in terms of your risk model. Like in doing threat modeling, you are thinking about external factors that may be interacting with it. So definitely tap on the system engineering people within your organization to -- if you yourselves are not steeped in system engineering methodology, but that's absolutely critical.

MS. ROSSI: So I'll follow up on that and actually reiterate something that was said on the last panel, bringing in your trusted partners in the industry to threat model the same for risk assessment. Looking across the industry, we cannot work in silos, right? We have to work together. So bring in those trusted partners, bring in those strategic collaborations that help us be better and continue to improve upon our product and our risk assessment, and I think that that's just a great best practice that we can continue to mature.

MR. ROTHSTEIN: So let's shift gears a little bit. That was the risk assessment, I would say, more on the, you know, design side. The other piece of the guidance that our panel has been tasked with talking about is the labeling piece, right? So the postmarket guidance has a significant section on labeling and how that relates to cybersecurity of the device. So a question for the panel is how are your organizations considering the end user when it comes to device labeling and, in particular, with a focus on the cybersecurity aspects?

MR. MORGAN: Well, Chris just -- for me, should we tell everybody to use the JSP if we want to, you know, align the harmonization? And in the JSP there's a section in there talking about customer security documentation which, in essence, aligns up with product labeling. It talks about the technical components, the software firewall configurations, your software bill of materials, cybersecurity bill of materials, whichever we want to call it, and really just all the security controls built into that solution, as well as the known risks that you have in your product, to share with customers. You know, there might be some gaps between what's in there and what's in the FDA's draft guidance but, in essence, customers are asking for it today, they want that information.

I think the piece that we have to figure out from the labeling standpoint is if we're able to provide routine security updates and patches to devices in the field, does that inherently conflict with product labeling requirements in the guidance? Does that mean we have to go back through and notify the FDA of labeling changes because we might have made a security update? So in terms of, you know, the guidance, Seth, that's a piece of feedback that, you know, we have from our end, too.

MR. FISCHER: I would like to add to this one, also coming again from the system point of view, and then speaking about the customer or the end user or the person reading this document. I think here is the major challenge, especially also now I think, for the guidance, what level of details you should provide in which form. I mentioned the system -- this modeling language, the block diagrams. I was tasked in my former company by yes, nice models, please make a picture there so that you can show it to management.

So I think everyone has to add in the issues if you're now starting with all this level of features, threat models, maybe risk analysis, maybe you explain how your scoring system is working, what was the model behind it. Maybe only experts can read this. And then again, in my area, if you're planning on speaking to healthcare professionals, maybe also in a hospital, their security specialists are available and maybe they can use that. But now thinking about an insulin pump with a patient, maybe the patient could not even read the manual because he needs glasses. And there are physician aspects -- so I think there will be a new set of challenges with this information and I think here, all of this is very important how we have to do this.

MS. ROSSI: So, from my perspective, the labeling piece is really where we bring it all together, right? We're talking about threat modeling, risk assessment. This is how we communicate this out, this is how we build upon that partnership and bring those essential pieces of information to our customers and we work together on solutions and bring awareness to risk mitigations, compensating controls, what have you.

Again, based on what Colin just said, I think I'll plug the Joint Security Plan one more time because it does provide a really great toolkit for here is security documentation, here are some of the things in everyday business language, not getting so technical that you're going over the heads of folks that might not understand it if it's written in language that maybe you'd have to be an engineer to understand.

It's a great concept for bringing to life a lot of the ideals that are in the guidance from FDA and we also, if I'll -- I'll say one more time, we've actually posted on our website, in that spirit of transparency, a copy of our BD product security white paper template so that any organization can come along and just take a look at what is it that we are producing for our customers, how do we take all the information based on the testing and the analysis that we are doing internally and making that available to our customers so that they know exactly, you know, what we are seeing and what we know and what we need to communicate out.

So whether you're looking at the JSP or you're looking at other examples, I think having that security documentation and providing that back is just, again, a really essential part of bringing all of this together, bringing it to life, because that's how we're going to communicate and build upon and continue to improve.

MR. HOYME: So, from a systems perspective, I think what we are touching is the recognition that there are additional stakeholders that have to be thought about at the point when you do the requirements for your system and the documentation. It has been traditional medical device companies that our sales and marketing people really know the physician and really will get the details about what they expect and what they need to use the device. But they don't even think about or necessarily have the connections with the IT and biomedical engineering groups that have to maintain and support it.

That said, there are different use environments for even those parts of the company and providing threat models, providing that information, would be great for those people in the customer complex that's doing risk management of the overall network. But then you've got the person who's going out to set up the device on the line and needs to know how to configure it, and if you hide the configuration data in reams of threat diagrams and that kind of stuff, you're going to get the classic too long, didn't read, and they're going to try to infer what they need to do to set it up by looking at the user interface and try to figure out what their choices are. So I think, as we think about documentation, you need to think about who's using it, when, for what purpose, so your goal is that when a device is being used, it's set into a secure configuration.

I'll then also reiterate what Colin said, which is there's labeling and there's documentation and the draft guidance has put things into labeling. The intent of the FDA is to try to minimize the resistance in the path between identifying a problem and getting a patch out and we certainly want to make sure that things like CBOM, which would change with every version of software, doesn't get -- doesn't get delayed because we have to go through labeling approval. So we want the intent without the delay.

DR. CARMODY: Just a quick note. I think people have been touching on it, but I just wanted to solidify the concept. The labeling and the section in the guidance is really intended to enable healthcare delivery organizations to manage their risk. That's it. So in terms of what that looks like, that's an open conversation, right? So CBOM is a perfect example of an open conversation that we need your help with so that when the guidance does go final that we can say, you know, here's the breadth and depth and frequency and all of the implications in terms of do we need to submit for a labeling change and all of that stuff. So that's very much an open conversation.

MR. BITZA: So the small thing I want to add to that, and it's somewhat building upon or restating what some of the other panelists have said, you have to think about the dialogue between the different stakeholders when your device leaves the manufacturer. What's in the back in the JSP tends to have a heavy clinical engineering focus. Don't forget about the dialogue that will take place between the physician and the patient. Is there information in there that allows the physician to intelligently discuss this with the patient so the patient understands the security life cycle impacts this particular product may have? So you're probably introducing some labeling you might not otherwise have thought of and it wouldn't be a bad idea to put it through some usability studies.

MR. ROTHSTEIN: One more comment from the panel, and then we'll go to an audience question.

MR. FISCHER: And I want to jump on this comment again. Thinking about the audience who has to read this, I heard a lot about people speaking about cybersecurity specialists and so on working on this topic, for sure, because they're here. But the point is, thinking now again about the patient, could he read this, could he understand this, what's he doing with this information. And, again, also you're asking maybe at some point in time we should think about how to -- this kind of information so that it could be used and not just a piece of paper because I already had interest from customers. Yes, we know that you have this -- but could you use this, because I've always thrown this away if it's a package, for example, for strips because I got it again and again, I don't need it. So there's also a question about what must be on paper, written, and what must, may be provided online. But when, if it's provided online, then the next question is okay, first thing I'm speaking about patient who has no internet access or even don't know how to use a PC.

MR. ROTHSTEIN: Okay, thanks. And we have time, I think, for one or two audience questions, and I see one is already up.

DR. VASSERMAN: Eugene Vasserman, Kansas State University.

The FDA persona of me is going to stay silent. The academic persona of me can't help but think, based on what the few people said of Shannon's Maxim, and that is a common security design principle which states simply the adversary knows the system.

And at least two people speaking for two companies have said we publish as much of this information as possible and I want to point out that that helps everybody. That helps, to some extent, regulators, that helps figure out what should be done if something goes wrong, that helps the company because the documentation must be polished, it even helps competitors because they can see what someone else is doing. It helps everyone, but it does not help the attacker because a correctly designed system, this information should be assumed to be known by the attacker already.

So I am really happy to have heard that. Publish as much as possible, maybe no one will read it, but it forces you to revisit what you're saying, it forces you to look at your own documentation and lead by example. You become the leader, and there's a competitive advantage of that.

MR. ROTHSTEIN: Any responses from the panel?

MR. MORGAN: I would just say we do get some customers that ask us to scrub some of the data. They don't sometimes want things like pen test results or vulnerability scan results. Some of that's a little more confidential, per se, and there's concern about if that gets leaked out. So sometimes some of the stuff gets pulled back, but yeah.

(Off microphone comment.)

MR. MORGAN: Yeah.

MR. ROTHSTEIN: Okay. Well, I got the time's up sign, so I apologize to the gentleman who had the question. So at this point, please help me in thanking the panel and I will ask, I guess, somebody from FDA to let us know what the next steps are with the workshop, since it is the lunch break.

(Applause.)

MR. D'AMICO: Thanks, everyone. Okay, so for the next steps, we're going to break for lunch, so if people want to stretch their legs, head to the restroom, and grab their food, that would be great. As Aftin said, this is going to be a working lunch, so please go and grab your food, and we're going to reconvene at 12:30.

(Whereupon, at 12:05 p.m., a lunch recess was taken.)

AFTERNOON SESSION

(12:30 p.m.)

DR. ROSS: Okay. Good afternoon, everyone. Thank you for coming back in after lunch and to enjoy your lunch while we have a keynote. Our keynote today is going to be from Ben Miller. Ben Miller is the vice president of threat operations at the industrial cybersecurity company Dragos, Inc., where he leads a team of analysts in performing active defense inside of ICS and SCADA networks. In this capacity he's responsible for performing threat hunting, incident response, and malware analysis mission for the industrial community, and he's going to share a few words with us today. Please give a warm welcome to Ben Miller.

(Applause.)

MR. MILLER: Hi, everybody. Thanks. Thanks to the FDA for inviting me, and thanks for all of you here, doing the work that you're doing. It's really important, really critical, especially as I age, like I want to have some sense of security as I go through life.

I guess, for my background, so I represent kind of three roles in how I'll be describing some of our -- some of the material I have. The first is, obviously, is Dragos. Dragos is a software company, we deploy out software that actively defends industrial control systems environments. A lot of the knowledge needed for that is not in a software developer's mind. So we have two separate teams at Dragos in addition to the platform: an intel team that's actively going out there looking for adversaries that are focused on industrial control systems and then my team, which really does the incident response, the engagements with the customer where we're very customer facing. And the goal there is not really to be a profit center for Dragos but to bring that knowledge back into the platform and instill it and incorporate it.

Prior to being at Dragos, I worked at NERC, North American Electric Reliability

Corporation. So NERC is a nonprofit institution, they've been around since the '60s, since the Northeast blackout of '69. In the early days they were an industry-led organization that created standards for the electric system. Now, these standards range from, like, vegetation management, like how -- what's the distance between the transmission lines and vegetation all the way to physical security of substations, generation plants, as well as cybersecurity. So I worked on the cybersecurity side, as you might imagine, focused not just on the regulatory matters but also non-regulatory areas such as information sharing and spreading knowledge and kind of being the liaison between government as well as industry.

And then prior to NERC, I was also an asset owner, so I worked at what was at the time Constellation Energy, who is a large holding company of generation, transmission, and distribution utilities.

So my background, I do not have a lot of exposure to medical devices. Really, my goal here with this keynote is to illustrate what we've done in the electric sector, specifically, and kind of how we've grown and changed over time as the threats evolved. And this is less industry specific and more general to industrial control systems. I think it also holds up fairly well in our arena today as well, where going from the bottom and working our way up, we have too few people for some of the challenges that we have today.

I was, 2 weeks ago, at the S4 security conference, which is one of the larger industrial control system conferences that are out there, it's an annual event in Miami Beach. There was a record number of attendees this year, 400, and I think 23 attendees to S4. There were representatives of -- I think it was 27 countries were represented in that conference. Those are really small numbers compared to the overarching security community. So Black Hat, which is another annual event that is held at -- or at Las Vegas every year in August, they had 17,000 attendees. So we're two orders of magnitude difference between what's going on in the critical infrastructure space and what's going on in the general, broader cybersecurity community. And that affects everything from how architecture is done, how regulations are understood and enforced, to how we're actively defending these systems, which really goes into the landscape of threats that are out there and what's really known. I'll talk through what we know about some of the threats on the electric side as well as critical infrastructure, in general, but essentially, they're anecdotes. We have a lot of stories that we've been able to piece together over the last 15 or so years, but we don't have concrete data or a broad understanding of all the threats that are out there.

So case in point, this is a good timeline of some of the activities, ranging from just complete lack of understanding or knowledge from the early days, going back about 20 years ago, to more recently kind of the impetus for a lot of security and critical infrastructure that's reached, kind of, buzzword status is Stuxnet, the first worm that was allegedly used or created by the U.S. Government in order to have impact to Iran's nuclear enriching capability. So that took out several of their enrichment products and basically damaged their equipment through a very stealth -- a virus or worm that was in their environment for months to years.

And then more recently, there's -- in 2013, 2015 there was espionage that was focused on malware that had a high degree of understanding of industrial control systems, so it's a recon tool specifically out there for these environments to better understand them and likely to be used as a portion of targeting, of understanding the attack sequence.

All the way to today. There had been two power blackouts in Ukraine, one in 2015, one in 2016, that the first one affected three different distribution companies. So when I say distribution, the neighborhood power company that handles, like, the poles that are in your neighborhood development, those are distribution companies. There were three that were impacted over the course of about an eight period -- an 8-hour time frame where they had blackouts. That was caused by adversaries directly manipulating their systems, actually using their systems to full effect by opening breakers and de-energizing power lines.

That was then codified in 2016 in another attack in Ukraine. This time instead of distribution companies, we focused -- the adversaries focused on the transmission, so transmission being the large steel structures, while you're on the highway, that are carrying high-voltage alternating current to various cities, those are transmission lines. In this case, a transmission substation north of Kiev was de-energized for approximately an hour. That was not caused by adversaries actively manipulating the system and so they codified that into a piece of software, malware, called Crash -- well, that we call Crash Override, that was specifically designed with various modules in it in order to affect technology that's only found in substations. So this is substation automation-specific protocols that were written and created inside of Crash Override, so this is a very tailored purpose towards that attack.

And then the most recent attack, in 2017 in Saudi Arabia, a refinery was -- the refinery tripped -- by trip of going offline purposely in order to prevent human loss, safety concerns. Those trips, two trips over the course of 2 months, were caused by a piece of malware called Trisis, also known as Triton, that would -- that was actually manipulating the safety controller, so this is the active safety system within the refinery to help protect the plant and protect the lives in it. That piece of malware was specifically targeting those safety controllers in order to have an impact. We don't believe the adversaries behind that were actually intending to trip those plants. We think it was some bugs in their malware. As they were testing it and deploying it, it caused those effects. So that's the threat landscape from the critical infrastructure sector in a nutshell.

But from an energy -- what was the energy sector doing over the course of that time frame? Really, in the late '90s or early 2000s, as an industry we were generally going through a modernization deregulation, so a deregulation creating of markets that helped set the price for electricity and being able to sell outside your territories and all of that -- in order of scale, when you're in an organization, if you're really a vertical organization that is regulated, distribution, transmission, generation are all within the same company. When you start de-regulating that, that means you have to start talking to other people, you need to create connections out. That's where technology really comes into play where the idea of well, we'll just keep that, that SCADA system controls our distribution area, we'll just keep that error gap.

Even if that was possible and that really did happen, which it didn't, it becomes -- it absolutely becomes impossible where you start asking the question well, how are you getting the seller information, how are you getting that over to the company that you're working with? You're not writing it down and then carry it over to another system and punching it through, that there is software connections, there's networks involved there that are creating that. So that, as well as the smart grid, which smart grid means lots of things, but essentially whether it's smart meters or storage and other technology, those all involve computer software that are running them.

From a regulatory body perspective, and where we are with regulation, so regulations -- I mentioned NERC as being -- originally an industry-led organization that created regulations that weren't enforceable, they weren't mandatory. In 2005 Congress approved the Federal Power Act which created a very confusing structure on how the electric grid is regulated, but essentially there is an organization under the Department of Energy called FERC, Federal Energy Reliability Commission, they are responsible for creating and improving -- well, approving standards, but the body that actually does the drafting and the enforcement of those standards is done by NERC, which is a nonprofit company that's representative of industry members that get together and receive their requirements from

FERC, the government, on what the goals of the regulations are. They go through a balloting process, a very long and strenuous sort of activity, to create, on the cybersecurity side, CIP standards, critical infrastructure protection standards. Those went into effect, the bulk of them, I want to say went into effect in 2009 and I was still an asset owner then, so I got the pleasures of implementing NERC's CIP protections within a utility. The challenge with those regulations, the challenge, maybe, with the electric sector in general is led and run by electric engineers and I joke, engineers are going to over-engineer things. And so from a regulatory aspect, the CIP standards are both very, very prescriptive and not prescriptive at the same time.

So there's a lot of expectations on how frequently activity happens, vulnerability assessments happen every year, for instance, but the -- how each entity -- each company enforces or creates their security program is very unique to themselves, how they protect their own information and do background checks. There are certain obligations from what a background check is, but how they monitor and enforce that is all tailored towards those individual companies.

And with that adds more complexity in that we're mapping regulations on how the grid works. The grid is essentially -- within North America, it's often cited as the largest machine in the world. If you think about it, both the U.S., portions of Mexico and Canada are a machine that's providing power. When I say machine, a turbine that is creating electricity that's being generated now, that's actually in sync with another turbine from Florida, it's in sync with a turbine that's in Maine. They're all in sync in order to produce electricity, so it's a very well large, orchestrated machine. And so the regulatory aspect of how we do that is tailored towards how the grid works.

There's four different interconnects, those are kind of independent bodies, the East Interconnect, the West Interconnect, and then Texas is always the Lone Star State, and then there's also a small portion in Nova Scotia, mostly due to its isolation geographically. So how the standards are in force are also broken up there. There's different teams doing audits in the west, in the east, in Texas. That means they have different standards and requirements of how they're doing the regulations as well.

I only have 5 minutes, so I'm going to jump to a couple different slides. One is going back to threats. So the Crash Override that I talked about, that's -- Dragos wouldn't call that a threat, really. We'd call that a piece of software. The threat is the activity group behind it, the adversaries that are doing the activity. We dubbed them as Electrum. There's another called Xenotime, which is fundamentally responsible for Trisis.

Out of the eight activity groups that Dragos tracks, those activity groups that are really focused on industrial control systems and doing a lot of reconnaissance and probing, only two of them have demonstrated ICS capabilities, that's Electrum and Xenotime. Electrum, really, forming out of 2016, so it's been about 3 years now since we've first seen them. Xenotime, over the last 2 years. And the others are catching up.

So the name of the game is that this isn't slowing down, this isn't necessarily getting better. The adversaries always have much more creativity, much more ability to adapt, while the defense is responding to the adversaries themselves.

One final point, I think, that's really helpful. I wasn't here for most of the morning, unfortunately, but I came in late. I did see there was a threat modeling session earlier. The challenge with threat in modeling, and I think this is applicable in your space as well as the electricity sector in general, is the threats that you're taking on, it's not just yours but it's also your customers. As you expand your customer base, the threats that are focused on those customers, you're inheriting their threats, you're inheriting their threat models, and you have to account for that in your technology and in your thought processes, how you pursue things.

So with that, I think I am, like, momentarily out of time. Do we have -- we do have some time for questions, though.

MR. TUGMAN: So first of all, hi. We don't know each other, but we are connected. I missed you at S4 but -- was there --

MR. MILLER: Oh, yeah.

MR. TUGMAN: -- teaching a panel there. So I wanted just to bring it home to the group, you mentioned CIP and CIP -- like, CIP 5 coming out in 2009, CIP 6 had to come in immediately thereafter. While it's a regulation and today we're talking about tiering of classes and protections --

MR. MILLER: Yeah.

MR. TUGMAN: -- but there's a relationship between what happened in CIP 5 to CIP 6 that they're talking about here today, with the classification of Class I to Class IIIs and that CIP 5, everybody, if you were BES Low, you didn't -- you were omitted from the regulation, the CIP 5 regulation.

MR. MILLER: Yeah.

MR. TUGMAN: They recognize that now everybody was CIP Low, right, so -- but could you talk about how they covered that with CIP 6 and then how they kind of had to go back and then -- to bring everybody back in the fold?

MR. MILLER: I will try my best. So I'm a bit long in the tooth on my CIP standards, quite frankly. Early on, like version one of the CIP standards, everything was risk-based approach, you're going to do a risk assessment and understand your risk, and then you're going to deploy the structure of CIP to that for those critical assets as it made sense. As they went through that process, it was very evident that the risk-based approach that the industries were using, each company got to define how they define their risk-based approach. It became evident during that phase that not everyone was not quite as -- maybe their risk-based approach was very helpful in the conclusions that they wanted to achieve, I would say, which removed assets and the protections that they had to do off of the table.

And then that's where the low, medium, highs came in and the question of, well, do I really need to treat my one substation that is practically a distribution-level substation the same as I would treat my control center, that's crazy. So then we started splitting those up into high, medium, and lows and having different protection levels there and basically, bright-line criteria that were defined in order to say if it's at this certain class of voltage then it's going to be a critical or a medium or a high and that's how they broke it down. The newer versions, quite frankly, with Version 6, I'm behind on it and --

(Off microphone comment.)

MR. MILLER: Got you, got you.

(Off microphone comment.)

MR. MILLER: Perfect. So I appreciate your time today. I hope to have provided some background on where the electricity sector kind of came from and from, not just the regulatory side but also the threats and the activities that we've seen that have been influential to just how we approach security. And quite frankly, I think the two industries that are similar is -- there's a lot of fear, uncertainty, doubt; a lot of very good headlines that can be created on both the electric sector side as well as the health side and that's something that yes, the threats are real and they're important, but they're not as serious as our imaginations can lead us to. So we have to address them but we have to be well measured in that as well. So thank you for your time. I appreciate it.

(Applause.)

DR. D'AMICO: Thank you again to our keynote, Ben Miller. Next we're going to move on to our fourth panel, Transitioning from Implied Trust to Trustworthiness: Authentication, Authorization, and Encryption.

(Pause.)

MR. HAZELETT: All right, welcome to the afternoon, and first off, I'm Matthew Hazelett. I'm a reviewer in the Implantable Electrophysiology Devices Branch, and this is our fourth panel session on transitioning from implied trust to trustworthiness including the authentication, authorization, and encryption design controls referenced in the guidance. So, first off, let's go ahead and introduce the panelists.

Tara.

MS. LARSON: Is this on, is it working? Hi, I'm Tara Larson. I work for Medtronic, and I kind of am the designer for the cardiac rhythm to heart failure disease management group.

MS. RICCI: Hi, I'm Linda Ricci. I'm the Associate Director in ODE focused on digital health, which includes cybersecurity implementation of policies.

MR. COLEY: I'm Steve Christey Coley with the MITRE Corporation supporting FDA on various efforts including cybersecurity analysis and vulnerability handling.

DR. VASSERMAN: I am still Eugene Vasserman from Kansas State University and a senior staff fellow at FDA.

MR. HOYME: I'm Ken Hoyme. I checked during the lunch break; nothing I said in the last panel got me fired, so I'm still with Boston Scientific.

(Laughter.)

MR. CHAPMAN: And I'm Joe Chapman. I'm a principal hardware security engineer at the MITRE Corporation, also supporting FDA.

MR. HAZELETT: All right. Well, so we've heard from a prior panel discussion to trust no one, so now we're tasked with discussing trustworthiness. So I think I'd like to start the -- get the conversation started around discussing implementing the design controls around establishing trust and how to go about ensuring that devices can be trusted in their day-today operations. So I wanted to kind of get the ball started. Maybe, Joe, if you had any thoughts on what the core basics of establishing trust in communication would be?

MR. CHAPMAN: Sure. So thank you for putting me on the spot first; appreciate that. So when it comes to trust and trustworthiness, I think those are two distinct concepts and I think it's important to get terminology kind of set for the panel for the ensuing discussion.

So trustworthiness, to me, it's a foundational property of a system or a protocol or an algorithm or some process that you're incorporating into your design. And so trustworthiness is, basically, it's really a complex thing to establish because there's a system architecture component involved. So if my system does not have features in it to support a cryptographic primitive such as, like, authentication, then it's difficult for it to be trustworthy because it has no way to prove it is what it says it is and wants to do what it says it wants to do.

But there's a lot of other components to trustworthiness, and I think we can start thinking about the bill of materials and the supply chains involved there. So if I incorporate a whole bunch of untrustworthy software into my design, when it comes to implementation time, am I still trustworthy even though the design looks good?

I think it's a bit about policies and procedures for the company involved, so if I don't have any kind of background check or internal security reviews and audits in my own company, can I produce trustworthy products? Can I be trusted to operate, you know, a server or in a cloud environment? So there's a lot of that that goes into trustworthiness, and I'd say that's kind of an upper bound for the amount of trust that we placed in a system before -- you know, before we have some sort of evidence that's been changed to proof.

So trust, trust is a belief at its core. It's a belief that something is not going to violate this assumption that I have. And so, you know, trust without trustworthiness is a dangerous situation, right? So I think that's the -- the point of this panel is designing systems that can be trustworthy and so on and so forth. So I think that's going to wrap up my 2 minutes on the spot, thanks. Open up for comments.

DR. VASSERMAN: The previous keynote mentioned power systems and separating things by voltages, that is high, medium, low, and critical being above high. If you have one -- so you have a number of assets within your system, they are all trusted to do something. The design challenge would be are they trustworthy to do the thing that you assume they're going to do, the thing that you trust them to do? A public key infrastructure, TLS, is only as trustworthy as its weakest link because every chain in that link is equally trusted. So we see that that may not be the best approach.

I'm not going to talk about this yet, but I want to point out there were at least three terms I heard that were undefined: trust, trustworthiness and establish; that is, what does it mean to establish, what is a threshold that we have to pass, and I'm making an assumption there is, in fact, a threshold.

The other thing to which I wanted to reply, Joe, that you said is crypto. The reason we need cryptography is because cryptography is the science of information assurance, it is the mathematics behind information assurance. If you're going to prove something and you're going to do that in a machine-to-machine way, you can take several approaches but if you don't trust the communication medium and you don't know who's in between those two machines, you have to do something with cryptography. By definition, it is the science of protecting data authentication of encryption, of authorization. Or the mathematics of.

MR. HOYME: I'll just plug in that the amount of trust you require is probably a function of the impact of that trust failing. So yeah. So the mechanisms you use, you know, why -- when is two-factor authentication required? Well, it depends on what's the impact of you letting somebody through that shouldn't. We had one level of trust for the people that made our sandwich today. I probably have a different level of trust for those who manage my 401(k). So part of that also relates to the levels that we expect from -- you know, based on what the impact is of our medical devices on patient safety or on information.

MR. COLEY: Your 401(k) isn't going to kill you; the sandwich might.

(Laughter.)

MR. COLEY: That was not a comment on the actual food.

MR. HOYME: No, no.

MR. COLEY: Just the potential.

MR. HOYME: No, but --

MR. COLEY: Which actually illustrates a point, I meant to do that.

MR. HOYME: No.

MR. COLEY: I'm a professor.

UNIDENTIFIED SPEAKER: The 401(k) does provide defense in death, though.

(Laughter.)

MR. COLEY: Yes. Even jokes here are a means to an end.

(Off microphone comment.)

MR. HAZELETT: I think one thing I wanted to pivot on that Ken mentioned was in establishing trust based off of the constraints of the environment. How do you go about assessing the level of trustworthiness required when you're in different use environments for devices? So we've focused a lot today on HDOs, how does that change with the home healthcare environment when you're using patient cell phones as a conduit to communicate with devices when you have devices at home versus in a hospital.

DR. VASSERMAN: I want to maybe bring up end-to-end trust and then toss the ball to the other side to see if people run with that. Oh, that was it.

(Laughter.)

MS. LARSON: So I think that's a great question. So I think when you're designing the different systems you have to consider the environment that they'll then be used. Patients want mobile applications, patients want that control, they want to be able to have a smartphone or a smart device that is not obvious to everybody around them but with that you have to come with some kind of -- you have to commit some level of trust because you can't trust that platform. So the applications have to build that in, in a way that makes them independent of the risks of the smart device that the patient is using and in that way then you can actually trust the patient to know that they're taking an action and you, as the manufacturer, have actually controlled that risk for them. And then just kind of disclose that to them along the way in the design.

MS. RICCI: Yeah, I think it's really important to understand the needs of both the environment and the people that will be using these devices. You know, just to echo what Tara said, we need to develop devices that are useful by the people that they're intended to be used by and, you know, if that means that we need mobile platforms because that gets us the effectiveness of the device, then we need to be able to secure the applications on those devices. It goes to the balancing act of making sure that you are designing a device for the -- for its intention and making sure that it's trustworthy for its intended use and not designing just a black box that is a hundred percent secure that no one will use. So making sure we have that balance, I think, is very important.

MR. COLEY: One thing that I think is important to consider when talking about and thinking about trust is what are the assumptions that are being made as part of the decision to trust? In some cases, and I think probably with many, many patients and perhaps many HDOs, there is an assumption upon point of purchase and deployment that the device is going to be safe for use and protected from malicious parties. But, you know, that's an assumption that simply in its early -- isn't deserved until you really investigate more closely what the particular assumptions are that are made in how the device is designed and implemented. Where security vulnerabilities can come into play is when those assumptions that are being made can actually be violated. Assumptions can be implicit or explicit.

MR. HOYME: I think it's also important to know -- and I think this is a challenge in the guidance as it's currently written, there will be feedback -- that a device can have different trust relationships with different entities simultaneously. The model I think of is we talk a lot about equipment in a surgical room and what's appropriate during surgery when people are scrubbed up, about what the user interface should do and is the equipment acting on behalf of the surgeon with surgeon user authentication controls? If it's hooked to an EHR and pulling imagery data up to help guide the surgery, is that on behalf of the user controls or do we -- you know, the guidance talks about other devices and I think as we get to device-device connections we have to recognize that the trustworthiness of another device has similar needs, as a device, to a user, which is a human, and so therefore, again, the trust relationship has to be commensurate with the risks involved with what that device/human may be able to do with that information.

MS. LARSON: And to build on it, Ken just said we also have to consider -- when we consider the theory that the device is working and we also have to look at the use cases and then trusting your physician versus trusting somebody at your home or trusting someone in a subway station, just looking at the different actions that people can take around you and design those controls based on those use cases. I think that's going to be a lot of the assumptions that we're making in this case is who's around and what's actually happening.

DR. VASSERMAN: I just want to put Ken on the spot here. So are you saying we should or should not differentiate trust in (a) different use environments and (b) between different entities?

MR. HOYME: We should.

DR. VASSERMAN: Okay. Yes.

MR. HOYME: We need to.

DR. VASSERMAN: I'll take that one step further. I would say you can change the trustworthiness or the trust given to a device based on the context of use, for example, a locked OR with a guard in front of it -- I don't know how many HDOs do that. All of them, right? Versus a home-use device or versus a device that's barely attended because no one's really watching it.

But I also want to point out that there may be different amounts of trust that a device may put into another device based on the criticality of the output of that device, the function of the other device. I really should've named them.

Again, the level of trust should -- or you should look at trust based on how critical something is. If your temperature sensor is absolutely critical to the function of your system but it's also being used by a thermostat in the room, the thermostat doesn't really care much but the critical component that's using it for treatment must place significant trust. It's actually the same device that's being trusted but at very different levels based on what the -- for what the data is being used. That's good.

MS. LARSON: Well, you can just build it out, wouldn't you, at that point and want to build in some kind of whitelisting to make that trust implied and allow some other communications to actually communicate with that critical piece and that's exactly what we're looking for, is ensuring that you know what is communicating and managing that appropriately.

MR. COLEY: Yeah, at some point and at some level of detail you have to assign trust to some of your lower-level components. This is where a lot of the supply chain stuff comes in to. You're probably not going to implement your own operating system for your own, you know, thermometers or something along those lines, and even if you do, who's to know that you did it -- you did it correctly. So you get to a particular point in your analysis and in terms of deciding where am I just fundamentally trusting the correctness of the behavior of some of these components and how confident am I that they are reasonably trustworthy.

DR. VASSERMAN: That's a really good point. Especially in the security world I would trust a well-known and well-tested security library far more than I would trust something I built myself, I know I'm unqualified to build one.

MR. CHAPMAN: Yeah. I just wanted to take a moment just to talk about implicit trust. We've been talking a lot about these topics and for folks who might not even understand what we mean when we're talking implied trust, you know, as the title of the panel is up here on the screens, to me, basically, I would start by looking at interfaces in your system, so where are --

DR. VASSERMAN: What's an interface?

MR. CHAPMAN: Where are the interfaces -- sorry?

DR. VASSERMAN: What's an interface?

MR. CHAPMAN: Oh, boy. A boundary --

DR. VASSERMAN: Sorry, I'm not being --

MR. CHAPMAN: -- some sort of boundary.

DR. VASSERMAN: I'm not being difficult. I genuinely want --

(Off microphone comment.)

DR. VASSERMAN: Yeah, I genuinely want to make sure I'm on the same page, and as Ken points out, I'm being difficult.

MR. CHAPMAN: I would define it as a boundary between two entities attempting to communicate with each other in some manner or form and that's a very -- I'm trying to be as abstract as possible, so that can be people to machines, that can be machines to machines, it can be processes to processes, that can be threads to threads, right, and so on and so forth. So anywhere there's a logical boundary between two endpoints trying to communicate, that, to me, is an interface.

DR. VASSERMAN: You pass.

MR. CHAPMAN: Okay. Thank you, Professor. So, anyway, if you don't have any kind of a challenge or a sponsor or evidence being yielded or some other way of proving the sort of veracity and the authenticity of the claim, you're implicitly trusting the other endpoint in that interface and so simply relying -- in a system design context, simply relying on the fact that someone's able to communicate on an interface properly is what I think we're talking about when we talk about implied trust. That's not necessarily good enough in many contexts, perhaps some it is, but in many contexts that's probably not good enough. And so that's the risk of implied trust, and I think it's important to identify those scenarios in your systems, and I think it's important to either at least identify them and communicate them if not mitigate.

DR. VASSERMAN: I'm not sure whether I'm drawing the correct example, but maybe this is the argument of no one knows what wireless protocol we're using or no one can construct the same cable we're using to connect to our device, not to name anyone.

MR. HOYME: Yeah, was that it? I just want to make one more -- again, in anything in this kind of space, it all depends. As part of understanding the impact of putting a control in place to achieve higher trust you need to understand, from a usability perspective and what environment, what it is. For example, if you require a device, through its user interface, to challenge somebody with user credentials, you will either have something unique to that device, which everyone in a hospital hates because now I have another username/password or you have to have the device synchronized to active directory or some other kind of directory structure which has a policy that, because of various different types of rules, I have to change your password, it has to be complicated, and now if you have a device that is purely in a surgery room where you know you have a level of physical control about who has access, is requiring the surgeon to log in before they scrub in with an interface, there can be usability balances issues to say I'm going to get this additional trust but/or can I count on some other aspect of what allows the environment to be trusted which may be completely different from what credentials or things might be used to connect that device on a network and perform functions over that interface.

DR. VASSERMAN: An example might be these unknown future uses, you design something to be used in one way then suddenly someone wants to use it for telemedicine and sees nothing problematic with it without examining the internals of the functionality of the device because they didn't build it. So they connect it to a network because it has a network port and they connect it to an external network and then they connect to it. The problem is so can everyone else.

MS. RICCI: So I think this brings up an interesting point about when you are trusting something because of the environment that it's in or assigning trustworthiness to it because of the environment that it's in. I think it's important to actually document that and document the implications of that trust such that in the future when things change -- I mean, nothing really ever changes, right? So in the future, when things change, it can be reviewed and it can be understood and I think without having those -- that documentation, that word I'm searching for and can't come up with, to demonstrate why you're doing that, then you don't -- when you go -- the next time you go to build this you won't know why you made that decision, so coming up with the appropriate rationale for why you think what you have done is adequate given the environment.

MS. LARSON: And in addition, on top of all of that, it's important to understand that we already know that these devices will be used in cases we hadn't already planned for, so build that foundation in during the threat modeling phase and really understand those abuse cases and start to build that extensibility in. And one of the biggest problems we face is that long-term extensibility, and if you're not thinking about the lifespan of the device or the many use cases potentially available to it along the term of its lifespan, you will continue to have these problems with design for security and authorization and trustworthiness, and all that has to be done as part of the design level and thought about up front.

MR. HAZELETT: Sir, I want to get to your question.

UNIDENTIFIED SPEAKER: Sure. So nowadays we're building medical devices on top of lots of operating systems, you have Microsoft, you have Android-based devices like our own, Apple, whatnot. But there are also lots of snippets of other small bits of software that you may include into your overall device. Now, one might associate a lot more trustworthiness, maybe, or trust, whatever, to the operating system builder because they have more resources, they have a larger user base versus a small piece of software that you've now incorporated.

There's a little bit of a conflict, at least we see, where you say if risk is based on -- or the mitigations are based on risk; however, on the other hand you say that you're only as good as the weakest link in your chain, that sort of falls flat on the risk-based approach. So how do we -- how does the guidance -- or does the guidance propose to provide some clarity on that, because we face that issue on a day-to-day basis. We're not experts in the operating system itself, although we can get educated rather fast, but we don't -- you know, it's not reasonable to go through 15,000 lists of bugs that come out every 3 months from the operating system manufacturer. So does the panel have an opinion there on that dichotomy between the weakest link versus a risk-based approach?

DR. VASSERMAN: I can't speak to the guidance, I can't say what the guidance intended, but, as I said, for example, I'm not qualified to build cryptographic software and

I'm not qualified to build an operating system. So just as a pure example, I would look at the history of that particular operating system, but I would also ask myself a question in terms of my attack surface: Do I need the entire operating system? Is there an embedded version of it, a low-power version, a low-service version of it that I can use, which will still be good for my device but does not have a lot of services that I don't need, thus potentially exposing myself?

So the trustworthiness of the operating system is maybe not directly but certainly proportional to the services that it offers, to the number of services that it offers and the size of the operating system. At least for that one aspect of the question is what I can say. Of course, it is a weak link. I don't necessarily automatically think it's the weakest link, but other than the attack surface, I would also ask is it still supported and is the assumption that it's going to be patched for the lifetime of the device consistent with the stated lifetime of the device?

MR. CHAPMAN: I would -- sorry, okay. So I'd like to -- I'd just like to comment, but again, I don't speak for the FDA and I'm not the FDA, but my personal opinion on this is that risk, it's always about risk, right? We're trying to protect people, we're trying to protect the safety of people and so for me, I think a risk-based argument trumps other arguments; however, I think that the -- you know, the chain analysis and sort of getting a good view of what's in your system and where the critical components of your system are, I feel like that process itself helps you go and identify where you need to do more research and really scrutinize the software that you're building into your system. So not all software necessarily is going to be equal, but at least going through kind of the risk-based approach and doing the threat modeling exercise we discussed earlier helps you identify those key unknowns in your system, your key, you know, gaps of understanding at the current time. That way you can at least call that out as areas for future research for your R&D teams and your development teams.

DR. VASSERMAN: This is, of course, a direct tie-in to CBOM. I actually like SecBOM, but that's not my choice.

MS. RICCI: Yeah, I was going to build on what my fellow panelists have talked about. It's about understanding the implications of what you're putting into your system. So, you know, the attack surface and the threat model should document, you know, what you think are the weak points in your system. You know, if that includes third-party software that you don't have a control over, you should understand what the implications of something happening to that are. You should understand what you're going to do, as a manufacturer, if something comes up in that software. So you really need to, you know, in addition to just noting it and saying, well, I have to trust this because I don't really have any other choice, understand what the implications of that trust are and make sure you can act on them.

MS. LARSON: And to build on what Linda said, you can use that risk-based approach to see if there's -- if the risk is too high, you can also design that software component out, or you can make another design choice, or you can build a defense in depth layer, understanding that a risk-based approach and that risk of that software is what helps you drive that more intelligent design. You have to use third-party software, you have to use off-the-shelf software, but you also have to understand how you're designing that and how you're using it.

MR. HOYME: I'll make just a very brief add-on to it. Complexity is the bane of good cybersecurity, so a really, really well used, highly complicated operating system will have a very actionable stream of vulnerabilities being fixed constantly. A smaller streamlined operating system like one that has mathematical proofs might be -- we haven't seen it a lot in this industry, but we do see it in aviation and some other industries, that's something that's far more streamlined and has a smaller attack surface.

MR. COLEY: Just one point on top of what Ken said and to tie some of these previous points together. Unfortunately, a lot of this is easier said than done. Code does not behave as expected. Even a single third-party component may consist of millions of lines of code, each of which, if used in a slightly unexpected way, may introduce a vulnerability.

I forget who mentioned it earlier today, but they talked about security being an emergent property of a larger system and that's what happens when you have different pieces of code interacting in ways where they might work correctly together -- they might work correctly independently, but when you put them together you're effectively creating a new layer of code or something along those lines that introduces new kinds of interactions for which you're not necessarily prepared or anticipated.

So I think there are some significant challenges in terms of really understanding, down to a very low level, what all of the different kinds of interfaces may be and how things will interact in ways that might introduce security issues. Just in case anyone was thinking this problem would be solved tomorrow.

DR. VASSERMAN: That's a very good point. You have to look at the layers individually, but you also have to look at what you built as a system. By the way, by definition, this is all easier said than done.

MS. RICCI: So we'll give you until Friday, is that fair?

MR. HAZELETT: One additional question I came up with based off of the discussion, so there's been a lot of mention of implementing some third-party software solutions into the overall design of the system. So from an actual design perspective, again, you're looking from, as an initial state of trustworthiness, but how do you kind of foresee approaching as some of those solutions may no longer be supported for the life cycle of the device, tackling those issues that could present themselves down the line if you're relying on another company or software solution?

DR. VASSERMAN: This is another point that ties very well with what Joe said about interfaces. If you have a very good understanding of the interface between either the various layers or between that component, then you should be able to -- again, easier said than done, but swap it for another component that does as -- does something as close as possible to the original component.

But if the interface is not understood, if the interactions include side effects, then basically you have spaghetti code and it's essentially impossible. So I see this as a fantastic argument for a modularized design where you understand -- where a module does an individual function or several functions, you understand very well what it does, whose data it consumes, who consumes its data and what the format is and so forth, and don't forget to convert from imperial to metric.

MS. LARSON: Yeah, I agree with Eugene on the modular piece. Having been impacted by software that's no longer made or manufactured, it becomes an area where you have to have that modularity built in and be able to manage that, on top of that. Having more use of off-the-shelf software makes it more likely that you will never run into that problem of software not being supported by a manufacturer or a developer of it, helping to make sure that life cycle or patching is available to you and having a modular-based approach makes it so patching is not as difficult as it can be sometimes.

MR. HAZELETT: We have a couple questions lining up, so I'm going to try to get through as many as we can.

UNIDENTIFIED SPEAKER: Thank you, good afternoon. Specifically to the guidance that's coming forward and the conversation around cryptography being the underpinning of managing some of that trust, it would be nice to have inside the guidance the fact that cryptography itself has a shelf life and goes away, and be able to manage cryptography on devices that are in people's chests with a high degree of assurance to be able to replicate that out and to do it in scale over a common medium, an iPhone or something else that actually has to do the transmission of the device is something that, I believe, is critical going forward. Thank you.

MR. CHAPMAN: I'd just like to comment really quick. I completely agree with that, with that statement, thank you for making that point. There's one thing, you know, as we talk about trustworthiness of a system, right, the trustworthiness is a property of the system, Eugene asked an interesting question to me at the -- before the panel, actually, is trustworthiness itself immutable, in other words is it a static property? And I'd say no, it's actually a function of the world, it's a function of advances in analysis techniques, for example.

So, you know, classically, I think of it as RSA and quantum computing, right? So Shor's algorithm is going to break -- or I'd say crypto systems isn't as -- okay. It's a bit of a gray area, but it's presenting challenges for RSA-based crypto into the future. So I completely agree, it's got a shelf life that does need to be thought about in the design which is yet another challenge to put on, but I think it's a very good market. Thank you.

DR. VASSERMAN: Just to be slightly pedantic, I agree with the concept but I just want to point out that cryptography in and of itself doesn't have a shelf life. Algorithms and algorithm parameters have the shelf life.

MR. HAZELETT: Next question.

UNIDENTIFIED SPEAKER: So with the growth of home medicine and wearables, how do you deal with, perhaps, trustworthy software in a not trustworthy environment or maybe not even on trustworthy hardware?

MS. LARSON: So I think that's another area where modular comes in is you have to focus at the application level and assume that you can't trust the operating system, so build those controls into your application using the whitelisting manifest, whatever it takes, just assume that that OS is not safe and secure.

MR. HOYME: And, certainly, when you come to your availability arguments, you have to -- you know, there's interesting challenges about why you can do trustworthy things within the execution platform that you're assigned but you may not be able to be assured that you get resources that you need. So it just becomes a factor in what you can do and what you can trust.

MR. HAZELETT: Next question. Real quick.

MR. FERNANDO: So if we -- sorry, Anura Fernando from UL. So if we think about sort of the history of how the FDA guidances have evolved and so forth and look back to the late 1990s, we saw the early guidance on software validation and those types of things. If we now think about software as having more parameters to be considered from a cybersecurity point of view, I guess I have a two-part question.

First is what are your thoughts on leveraging some of the existing mechanisms, you know, that have emerged with the recognized standards recognized by FDA like IC 6304 for software development life cycle and 14971 for risk management and things like that, and the efforts to sort of extend those concepts down to the realm of cybersecurity and the role of standards in building a trust model because, you know, if you think about electrical safety, it's really the ongoing evolution of electrical safety standards that now prevents us from having to sit there and stare at the wall when we plug a device into the wall watching to see if either the device or the wall is going to catch on fire. Do you think we can get to that same point with security considering that it has to evolve? And what do you think will be the role of standards in that?

MR. HOYME: I think we can get to the point that cybersecurity causes our walls to start on fire.

(Laughter.)

MR. FERNANDO: I think we're there.

MR. HOYME: I'm pretty sure we're there, but having said that, perhaps running a provable OS and running a small piece of provable code on it you may be certain that it's going to last you a while unless some of the explicitly stated assumptions within those algorithms are violated or the hardware breaks.

MR. COLEY: I think from, sort of from a code perspective, software perspective, the tools that we currently have accessible to us sort of as an industry of doing code analysis have a number of different limitations that -- some of which may be able to be overcome, you know, academia, folks such as yourselves, working on improving the utility of tools, the ability to, you know, reduce the number of false positives, increase the number of true negatives. But code is so complex, I think there are just fundamentally certain kinds of limitations with respect to tooling and any other kinds of analyses that we're talking about let alone the next layer of complexity that's involved when you start plugging things together.

We have limitations in our tools in and of themselves, which is okay, that's the way the fact of life is, but I would like to get to and I keep hoping we can get to a point of, at the very least, understanding in some reasonably quantifiable fashion what the real limits of those tools are so we then can understand what the limits of our own assessments are.

MR. HOYME: So a good friend of mine coined a term that a fool with a tool is still a fool and that is one of the challenges, I think, which we haven't gotten to, I don't know if we have a panel discussion on this, which is the overall aspect of education in the workforce. I mean, in the end, there is a certain amount of automation you can do but you still need to be able to design secure systems and train people and get your -- the key people making decisions to understand how to make good decisions. So there's full employment in this space.

DR. VASSERMAN: I must speak very quickly. A very quick point. Software complexity is exponentially greater than any other engineered structure we have constructed. We must tread carefully.

MR. HAZELETT: All right, they're playing -- thank you, all.

(Applause.)

DR. D'AMICO: We'll head into our last panel of the day, Increasing Transparency, Advancing Protection, and Enabling Timely Response: Cybersecurity Bill of Materials.

(Pause.)

MS. RICCI: All right. I think we're ready to get started. This panel is about increasing transparency and enabling proactive action in the cybersecurity bill of materials or Eugene's security bill of materials, if you prefer. Before we get started on this panel, quickly, I want all the panelists to introduce themselves and then we're going to start off with two very quick overviews of the CBOM to help us form our panel discussions. So why don't we start down at this end down here?

MR. HORNBERGER: Zack Hornberger, Director of Cybersecurity at Medical Imaging Technology Labs.

MR. McNEIL: Michael McNeil, Global Products Security and Service Officer for Royal Philips.

MR. JACOBSON: Jim Jacobson, Chief of Product and Solution Security Officer for Siemens Healthineers.

MS. JUMP: Michelle Jump, Vice President of Cyber Program Initiatives at Nova Leah.

MS. RICCI: Linda Ricci, same job I had before.

MR. CORMAN: Josh Corman, founder of I Am The Cavalry and CSO for PTC.

MR. FRIEDMAN: Allan Friedman, Director of Cybersecurity Initiatives at NTIA in the U.S. Department of Commerce.

MR. ASKE: Jennings Aske, Chief Information Security Officer, New York-Presbyterian Hospital.

MR. ZALEVSKY: I'm Ken Zalevsky. I'm the head of medical device cybersecurity at Bayer.

MS. RICCI: All right, now I'd like to turn the microphone over to Allan to give a few opening remarks.

MR. FRIEDMAN: So first I'd like to thank Suzanne and Seth and Aftin and the rest of the FDA team for moving so quickly to realize that the government was now open again and so therefore I could come all the way up to White Oak and join this amazing conference, so thank you for the late inclusion.

One thing I want to say is that this notion of software transparency or the bill of materials is very much not new. A lot of the pioneers who have helped bring it to the forefront are in the room right now, are up on this panel, are sitting right next to me, so there's been a lot of great work that's happened here. And I also want to offer kudos to the FDA and, indeed, this entire community, of basically being drivers in advancing this notion that transparency can enable a lot of great work. In fact, we heard on the last panel, when you have transparency about third-party code you get a much greater understanding about trust.

NTIA moved into this space a couple of years ago because we see this as part of a broader conversation. This cannot be something that is purely a medical device issue for a number of reasons. One, there are a bunch of other people that have noticed gosh, transparency about third-party code could be really useful. Some of them have even larger checkbooks than this community. So DoD has said this might be something that we're going to care about. Commerce likes it when markets work and so when we start hearing about particular sector-specific solutions we get a little worried, especially when we are all facing very similar risks and derive similar benefits from transparency. And so the NTIA initiative around coming up with a shared solution for a software bill of materials was built to find common ground on a minimal viable product, how do we get started with what we have today, and also understanding what do we need, what doesn't exist, that we can start working on for the next generation.

And I want to close with two very important reasons of why the solution for a bill of materials and transparency has to span multiple sectors, it has to span the entire ecosystem. First is, of course, as I mentioned, no one benefits from a sector-specific solution. A lot of your organizations don't just make medical devices, you make other things. It would be really annoying to have to have multiple engineering products. And also, by the way, hospitals and HDOs and everyone else acquires IT from different strains, and so we want to have it really as helpful if the switches and the blinking boxes have similar technology.

But for a practical perspective, since this requires widespread adoption, not just among the industry players but among the open source components upon which all this is going to build, we really only get one chance as a community to try to push upstream the changes that we need to see from the open source community. And I know Jessica Wilkerson (ph.) from the Linux Foundation is here, and there are some other folks that have worked a lot with the open source community. We need to be able to present a fairly common vision of what we need to secure the ecosystem in the proprietary software domain to make sure that we can get what we need from the community that's making the underlying building blocks on which we all depend.

So, summarizing, we desperately need leadership and I want to thank the panel and everyone who's participating in the NTIA initiative for that leadership, it's not too late to join, come find me after, but we need to keep in mind that this has to reflect the broader solution set that all of us are going to use across the digital ecosystem.

MR. CORMAN: All right, I have 5 minutes to say 3 hours worth of things. This is a graphic that came out to the NTIA working group. Again, I'm Josh Corman. I've been working on the idea of software supply chain hygiene and bill of materials for over 6 years now, a lot of it liberally stolen from Deming and Toyota supply chain in the forties, so these are proven supply chain principles attempting to be applied to modern software development. The head and shoulders leader for adoption and maturation is in financial services. In fact, the largest software manufacturer in the world is a bank, more than Apple, Google, Amazon combined.

So when you look at these -- the software adoption, the idea from Deming is three principles which we outline as S1, 2, and 3, that you should use fewer and better suppliers of parts. Number two, you should use the highest quality supply from those high-quality suppliers. And number three is you should track which parts go where throughout the life cycle and retirement of goods so you can do a prompt and agile recall. So pick fewer and better suppliers of airbags, don't use a known vulnerable batch of airbags and if there is a bad batch, because things happen, do a prompt and agile recall based on the tracking.

So this kind of concept is most acutely adopted here in healthcare. Michael McNeil, myself, and some others had the privilege to serve on the congressional task force for healthcare cybersecurity. A few days before we started Hollywood Presbyterian shut down patient care for a week and diverted ambulances to other facilities. The root cause of that shutdown was a single Java de-serialization flaw in a single Java library that they were warned about in a single device. They were warned but they couldn't answer two simple questions: Am I affected and where am I affected? It was opaque; they were blind. It was an avoidable harm. So one of our strongest task force recommendations was to require a software bill of materials in all medical technologies. Congressional oversight through

House Energy and Commerce liked it and asked FDA and HHS to do it, it's now being done, and with Allan's help through the NTIA process it won't be narrowly designed just for medical, it will be something hopefully the whole ecosystem can use. So when people say it can't be done, it's being done. It's being done mostly for productivity boosts and profitability enablement in financial services.

So what we drew here, out of our working group, just to orient you, is we're going to talk today mostly about the premarket guidance for final good assemblers. Maybe that's that infusion pump. But that bedside infusion pump is going to be deployed through several hospitals who don't know if they're impacted and they don't write it from whole cloth, they take a bunch of supply chain, some of these things are going to go all the way to the beginning.

There's atomic individual parts, like a Log4j or a Bouncy Castle cryptography library, then they go into these big compound parts, like maybe an Apache Struts mega project or jQuery or something or a JBoss, which is what hit Hollywood Presbyterian. Some of those get aggregated together directly and indirectly and they make it into maybe a Michael McNeil Philips medical device, they're going to add their own special sauce and then ultimately that final good assembled will have an SBOM, that's the aggregate of lots and lots of little less SBOMs, but every part with every version could be passed downstream.

So, currently, when there's a flaw in maybe Apache Struts, you might be able to answer am I affected and where am I affected, were that shared, but you may not be able to. And folks like Kevin from earlier today, at Mayo Clinic, may be asking for these contractually, but when they get them they get one flavor from one vendor, a different flavor from another vendor, they're not machine readable, they're not compatible. So what we want to do is take this good practice of sharing the ingredients and the versions such that you can answer am I affected and where am I affected. What we want to do is harmonize the outputs of these tools so they look similar, machine readable, and can be done without a lot of human cutting and pasting. But number two, we want the FDA, in this case and others later, maybe DFARS and the federal Defense Department regulations, to kind of ask for these to increase adoptions. Number one is harmonize what we have available. Number two is amplify adoption. And number three is there's some amazing use cases we can't yet do without some new extension and innovation and we're trying to do that.

In the last minute I've been given here, it's really important to understand the chaining here. We don't want to look at patient health of a single stakeholder in this environment, we want to look at the aggregate value unlocked if we have a consistent line of sight, because it's not enough if one of those compound parts, like an Apache Strut has a flaw or maybe it's in the product I produce, like an exceeder or a ThingWorx might go into most of your medical devices which might affect a hospital, which might affect a downstream patient.

We want it such that instead of there being many, many months of trust and hope that people communicate those flaws and fix those flaws, that could take a year and a half, in the meantime to exploitation for adversaries is being compressed down to days and weeks, the mean time to remediation for the good guys is basically months and years. So we're trying to compress that by increasing line of sight such that any little flaw, any little baby project, anyone in the chain, could be a real-time indicator that you have something to assess, patch, mitigate, take offline, etc. So I'm hoping to structure some of those things. The S1, 2, 3 is in procurement, maybe by the best manufacturers who -- based on who can produce an SBOM, who has the best hygiene. Number two for S2, maybe your go-live testing looks at the relative -- the least vulnerable version of the strategic vendor you've chosen to depend upon. And then for number three, it's the ongoing caring and feeding.

Maybe it's according to vulnerability disclosure program if you're a manufacturer, maybe it's your vulnerability management program at the hospital. But can you pay attention to emergent events and respond quickly? So those are the frameworks we wish to capture, dozens of other use cases. Thank you.

MS. RICCI: Thank you very much. I think I'd like to start off this panel with just a general overview question about -- particularly targeting the manufacturers to start off with. What do you think are the best uses of a CBOM and how are you challenged to actually implement something like this?

MR. JACOBSON: So the primary use case, from our standpoint, is looking at the HDO and saying well, they didn't have a risk management system that they need to effectively have information about the devices in. So what we need to do is provide information to our customers, to the HDOs, about what risks they have or what vulnerabilities they may have as a result of the components and integrate that into their system, into whatever process that they use for risk management. Whatever tools that they use, there has to be a way to get the information from one side to the other in a way that we can get some experience with to identify that it's being effective

So the work that we're doing with the NTIA in establishing a proof of concept for SBOM or CBOM, or whatever you want to call it, is critical because we keep talking about trust here during this workshop and one of the elements of trust that we need to establish is that manufacturers can provide information and it would be consumed effectively by our customers. So the actual process from the manufacturer's standpoint isn't particularly a challenge if you have -- if you have a vulnerability management process already in place, that is, a manufacturer, in order to manage the vulnerabilities in their products would have to have the software bill of materials already present within their system. So it's a matter of what process we use to expose that information and what -- a process that the HDO uses to consume that information and exercise the use cases that we've identified, use cases like for risk management, use cases like procurement, for instance.

MR. McNEIL: So in addition to what Jim had just stated, I mean, total agreement, again, it's the ability for the information that is disseminated transparently and how it is consumed. That's interesting when Josh kind of walked us through some of the NTIA models and activities that we're doing. That consumption piece has been, you know, the Achilles heel in this entire discussion. If you ask one consumer, health delivery organization, about what their needs are, I have yet to have anyone give us this schematic in terms of what their response would be and to the level or degree. What I will say is that what has been consistent is something that they can have from a machine readable and to be consumable has been a consistency.

And then from my own personal perspective, I've been extremely strong and a staunch advocate on the fact that as a manufacturer we have to be able to leverage, you know, one level up the communication. I'm not in a position to try to disseminate multiple different variations of the data and information to be able to support the marketplace, so we need to, as an ecosystem here, get consistent with the tools and the deliverables and how we would execute.

And, again, I'm in total alignment, I think, as Allan stated, because the SBOM, CBOM, whatever we want to describe it, I have to produce it, as well, for the toothbrushes and baby monitors and other solutions because I have one security by design process that we follow across at Philips, and I don't variate them based upon from a healthcare or the medical device side, just for comments.

DR. DASERIC: Yes, my name is Dr. Gene Daseric (ph.) from ICU Medical. Is on? Oh, I need to get closer. So okay, I am Dr. Gene Daseric from ICU Medical. I'm also an adjunct professor teaching computer science in a university in San Diego.

I understand trustworthiness, I truly understand transparency of -- especially with the CBOM, provide hospital enough information that they need in order to protect themselves from cyber attack. So the question of having -- CBOM, as you're releasing, if you provide it to the FDA and if it becomes public, do we have a way -- because as it is public, the bad guy would also see the same thing that the manufacturer can see, therefore they can take advantage of thing that is not an unknown of those products. Is there a way or do we have to create a way to prevent that information from getting outside to the bad guy in order to -- not to take advantage of weaknesses that we may not be -- I'm sorry -- we may not be aware of?

MR. FRIEDMAN: So I think the question is aren't we worried that if we have a list of ingredients, we have what if a third party came on and said the bad guys may also learn about this? And, first, I think it's important to draw a distinction between making this data available to the end user and making it public but at the same time we shouldn't be naive, if there is sharing of data, we must assume that there's a possibility that it will become public.

This came up at the very first NTIA meeting and it was roundly met with derision by the security research community because they can tell what's in your products today. The good guys can't. It's the bad guys who can with just a little effort to figure out what's under the hood and that's the advantage of providing schematics, right? Whether it's a car engine or any piece of hardware, anyone who knows what they're doing can open up the box and look. If you're trying to help people fix it or protect themselves, then you need to give them that same information.

DR. DASERIC: So on this info for the -- based on your information you're saying that at the beginning provide all the information to anybody that needs it, is that correct?

MR. FRIEDMAN: I think the exact nature of how we're going to -- what the active transparency is, is one of the questions that's being explored. So right now we'd be -- how is the data produced? That's what the MDMs are doing. How is the data going to be consumed, we need to understand that and also we need to understand what's the business case of all the great tools that are going to be built on it in the future.

There is a very real question of what does the active transparency look like. I think very few people are arguing today for complete publishing, somewhat, but I think the point is we shouldn't view that active transparency as introducing risk that didn't exist already. If someone wants to go after your product they can today and they probably are for competitive reasons, not just for malicious reasons.

MR. CORMAN: There's a wealth of videos that can be covered on this very question, and they're valuable to listen to because they're fair questions. One myth, just as a teaser to go watch more, is a lot of the licenses for open source already obligate you to declare them. We were looking at the open source licenses in our car rental yesterday on your iPhone, so they're already published, they're just not published with version information, which is material to vulnerability management. So we want to just enhance the defensive value.

MR. McNEIL: One of the comments that I just want to maybe close on, on the response for you is that currently today I would provide that information upon request and alignments directly to my customers and I have mechanisms in place so that the customers have the access and the entry to be able to ask and get that information. And then obviously, the second positioning that that is communicated is a part of the premarket submission information into the FDA. So when we do our risk assessments and our risk matrix information, we also make that same information available. Again, those, to me, at least as a medical device manufacturer, are two of the critical areas that where I see we need to make sure that that information is provided. And I think to your question gets more in alignment with how that information is potentially shared among the community within the ecosystem in a safe and transparent way.

MS. RICCI: I think this is a great segue into our next topic about what information and level of detail should actually be included in a CBOM, what is needed at each one of these levels to make this a cohesive process.

MS. JUMP: Yeah, I'd like to weigh in on that, actually, real quick because I think that as we imagine the development of the CBOM or SBOM process we need to see this in layers versus one initiation of a final product. In part because this isn't just a CBOM, this isn't just a list of materials, this is actually the whole process that is required to share it, to keep it updated, to manage it, and then to collate it all at the end point. So I think it's really important to think about this in getting something simple out before we get a really complicated list of a lot of different information.

I was at Archimedes last week, the conference, and talked to a number of hospital providers who knew I was part of NTIA's software transparency group and were begging us to get just a simple SBOM out so that they could start to know what's in their systems, they want the software and the version so that they can start to create that. Because if you think about that, if you start giving them some of this information, they can start building the infrastructure to receive it, right? And then if you get that basic road set down, you can start adding the additional information if they have the infrastructure. If you give them a lot of information they need to handle, it's much harder to set that up and get that running. So I think that there's a strong desire to get this implemented quickly and implementing it simply so they can start to at least know what they have and get that out to them, it will make a huge difference in the community and the industry overall.

MR. ZALEVSKY: Yeah, I'd like to support that position, Michelle. I hear the same thing at Bayer. We hear the same thing from hospitals. Nothing today and just a little bit of something, even simple information to pass on to them is very, very helpful. So yeah, I support that same position.

MR. HORNBERGER: And over at MITA we're working with a stakeholder group on the document that some of you may have heard of called the MDS2. A new version is currently being drafted. And I think I might use the term high impact rather than simple, that major version number and software name, as simple as it sounds really, when we've spoken to the stakeholders in that group is where they think the most effect will be had most quickly.

MR. JACOBSON: In the proof of concept that we're working with, with NTIA, for SBOM, we definitely -- as Michelle points out, we want to walk before we can run and the basic goal is to establish that we can communicate information, establish the basic communication about component -- identify the component, identify the version. There are other aspects of it that we're looking at that, looking at it from -- that could be opportunistically provided, things like vulnerability, information is not what we're looking at now because that's too dynamic at this point. But things like dependencies may be provided as well.

Identification is another aspect we're trying to tackle, how do we identify uniquely the component and that's also within scope of what we're doing. But our goal is to, as we said, walk before we can run, let's get out an initial trial of this and then leave the standards definitions groups to finalize a format that expresses what we've learned already in that proof of concept.

MS. RICCI: So in understanding that there's still a lot of work to be done in developing an appropriate level of detail and other information about how to actually make this a reality, what do you think are the most effective mechanisms for sharing this CBOM information? I mean, particularly, if we're looking at the chart there is a role for manufacturers to be both generators and consumers of this information, so -- and certainly healthcare organizations, the same way. So how do we, as an ecosystem, develop the right mechanisms for allowing this communication to flow so that we can all have the information we need?

MR. ASKE: I'll comment. So a couple things, actually, and I want to go back to some of the things that Josh was saying to tie this together and to respond to your question, but for me, when I think about healthcare and its information security narrative, it's a pretty core one. Historically, healthcare organizations have not invested what they need to, information security, we've not had dedicated leadership at hospitals or small practices and our vendors have basically sold us black boxes and we've not asked our vendors to -- you know, whether it's an EMR vendor or a medical device manufacturer, to actually be transparent. So that's changing, it's changing because of things like WannaCry and some of the large breaches that have happened, large fines. My institution's paid one.

But, really, to be mature from an InfoSec perspective, we need to have information, right? And so we need information that can help us in the procurement life cycle, that can help us with ongoing operations, vulnerability management. So as an example, let's say there's a very high-profile vulnerability like Heartbleed, for folks who remember that, that's a lovely vulnerability. I want to be able to really quickly go and search a database and say oh, what's possibly affected by this? So then I can engage the manufacturer, engage in defensive actions, like maybe taking a device off the network, other things like that. So healthcare organizations need this, this is part of us maturing as an industry as it relates to information security. We're going from not investing to actually now talking about something that you hear financial services is doing, so this is great.

One of the things that was happening early in the conversations around healthcare doing SBOM was that people were talking about developing a healthcare-specific SBOM, which is something that I think would be a really bad idea. Standards matter, we can look at standards for how we practice medicine, we can look at standards like html or ftp and when things are standardized, they tend to work and we don't need industry niche things. So tying this all together, one of the things that came out of the original NTIA event -- and that was in June, right?

UNIDENTIFIED SPEAKER: July.

MR. ASKE: July, okay. Basically, at the end of the day, Jim and I basically said we should do this, we should just pick a small number of devices, get a couple manufacturers, a couple of healthcare delivery organizations, let's publish this electronically, we're going to figure this out, we're going to learn lessons and it's going to basically show that this can be done, this is scalable.

We haven't finalized every aspect of how we're doing this, we're right now working on, you know, how we're going to pool the information into our CMDB, how my team will operationalize this for vulnerability management. Ken's helping out, we've got a bunch of great folks working on this, but the idea being that we can't let, you know, fear, uncertainty and doubt prevent us from moving forward, we've got to actually try this, learn from it and iterate, and ultimately do it in a way that's standardized so that that facilitates the transmission and the collection in the ecosystem that will be necessary because of all the third-party software components that basically are kind of bundled together in the stuff we use.

MR. McNEIL: One of the things I kind of build on what Jennings had just stated when he talked about his ability and what he wants to be able to do and how to react is number one, within -- inherent in that is the appropriate asset management. You need to know what solutions, what products, what information you need to be able to have that's in your environment. I don't think that Jennings needs to understand, you know, what's going on with Siemens or Philips if they're not in his environment and there's no threat that presents itself, so that's why going back to the form of communications and transparency. We need, as manufacturers, the ability to continue to leverage the processes that we currently have in order to be able to communicate, you know, with our customers so they can get access to that information.

And then as both Jim and Jennings had stated, we can build upon and build off of some of the existing communication frameworks, but test this along the way to make sure that we can leverage what comes out in overall standards that can be used not only in healthcare but across other industries as well.

MS. JUMP: Yeah, this is Michelle Jump again. I think that there's also the opportunity to realistically consider that there are software tools that you may want to consider for, kind of, serving as a vendor between a software vending machine, an SBOM vending machine, so that there can be some cross-communication centralized in one place versus everyone going one on one for each SBOM as well.

MR. GATES: Christopher Gates, Velentium. Two points I wanted to make here. One is, with an SBOM, medical device companies are extremely reluctant to ever come out and say our device is secure. I mean, even leaders like Philips don't come out and say this device is secure. They do all the background work, they do all the incredible amount of work and as a result of this, this frequently gets looked at as a cost sink for a company and something that we do, we have to do, but it doesn't really gain us a market advantage. Pay attention, manufacturers, this is where you can make yourself an advantage in the marketplace. If you publish an SBOM you're going to look better to the HDOs than the competitors who don't have one. If you keep that SBOM up to date and you're not running an old version of OpenSSL and therefore susceptible to Heartbleed, all right, you're going to look better to those customer bases, this is your chance to shine, okay? Take advantage of this. And you don't have to claim you're secure, all you have to do is publish an SBOM on your device.

Secondly, we're talking a lot about SBOMs here today and that's great. The guidance talks about a CBOM, because what they did was they introduced the concept of hardware into the SBOM. I have a lot of problems with that and I'd like to kick this back to the panel and see how they feel about hardware inclusions, especially in the light of things such as medical apps they're going to be running on a BYOD phone platform. So thank you.

MS. RICCI: Another great segue, thank you. So one of the questions that we wanted to tackle is what are the challenges of including hardware and software in a CBOM?

MR. CORMAN: At least from my NTIA working group, when we saw this, we kind of had a heart attack. We took a breath and a pause and we said well, wait, there's a good intent here, right, hardware can be bad. I think the trigger here, I'm speculating the trigger was Spectre and Meltdown.

The way we chose to handle it, at least at NTIA, and I was dying to see how the discussions happen today and tomorrow here, is what's possibly trackable within the current scope of Allan's great project is, is there an identifiable, discernible, unique firmware version running on some hardware? From my understanding on the task force, a lot of the device manufacturers will use a different network card or a different chipset, you know, whatever they grab off the pallet today. There isn't a lot of rigor there, but the other half of me says maybe we could encourage more rigor. So it's kind of my hope, this is an opinion, not a declaration, but it's my hope that we maybe incentivize the inclusion of hardware in a CBOM that if you have something that is more of a controlled issue, if there's a Spectre or Meltdown in the future and there's an easier pathway to regulatory -- what is it called, surveillance, right? Postmarket surveillance. It becomes easier if you've invested in a rigorous hardware component, but I think it would be very challenging. One of the nice things about the SBOM is despite the FUD being spread around, the practices exist and they are actually fairly mature, it's more about harmonization and adoption. There's not solved yet hardware problems in large parts of the government and other sectors are still trying to work on them academically.

MR. FRIEDMAN: Just very briefly, amplifying what Josh said. Hardware, I think, does give us some understanding of risk but I think at the first level -- one, say one of the particular risks we're worried about, right, there are things that if you don't have hardware support for, you're going to have a bad time, but my understanding is those are covered in a lot of the other areas of the premarket guidance. It is useful to know that it -- right. Software identification is still a very hard problem, I think we can solve it, but as Josh said, the hardware identification, a SKU does not give you the detail that you're looking for in the hardware level.

I'd encourage the FDA team to reach out to their government colleagues at DoD and most importantly, at Energy, where there is a lot of hardware-based issues in smart grid that they are still tackling and it's still very much at the research phase. I think there's a lot of progress we can make but at the moment, I don't think it's something that we can solve by bringing, you know, eight people together and say we've solved it for the entire ecosystem.

MR. McNEIL: And from a manufacturer's perspective, I think we've already discussed the fact that getting alignment and getting something that is transparent and consumable in an iterative process is a good approach and a way to go. I do think that we've also identified that even just on the "SBOM perspective" that there's complexity within that model and I do believe that adding the hardware component where we're at from a maturity perspective and from a development and a deployment stage, it really will make it a much more complex issue to execute.

MS. JUMP: Yeah, I completely agree with Michael. And I think the other issue here is SBOM is software bill of materials and bills of material are supposed to be component lists. And so I personally have an issue with the language of cybersecurity because you're not listing components of cybersecurity, you're listing software components or hardware components in this stage, so I think the language does matter because of the way bills of materials have always been used before. Manufacturers are familiar with this and I think we should stay that track, get this launched, focus on software because that's our biggest hit and then we can move on from there and really expand it to where we see it moving within the industry for risk-based decisions.

MR. JACOBSON: Yeah, and I'll just reemphasize that. Let's get the S part of it solved and then let the research continue, let the standardization efforts continue and include hardware as a goal, but we're not ready to do that as an organization. We're ready today to produce SBOMs, as an industry, to produce and consume it. So let's get it understood and then move on.

MS. RICCI: Thank you. Okay, we have about 1 minute left, so we're going to do speed questions. Go.

MR. TUGMAN: So really, you mentioned the smart grid. I actually chair a working group updating a DOE, updating the C2M2, and I'm also chairing a working group for the American Petroleum Institute, 1164 standard update, they're both dealing what that exact problem so it's awesome that you mention that. That is the actual problem, especially when you get into IIoT. So where I was going to go with that question before you kind of mentioned it was what's the outreach that's happening to the organizations to make sure, to Jennings' point, if you don't want to make an MDM-specific schema then there has to be an outreach to make sure that somebody else doesn't go into McNeil's point about toothbrushes, well, they're going to have an SBOM, CBOM, HBOM, whatever it is, so what is the collaboration or communication happening to make sure that we're not (1) reinventing the wheel and (2) that somebody doesn't invent over us.

MR. ASKE: Well, we're looking at SWID and SPDX, so using existing standards.

UNIDENTIFIED SPEAKER: Okay. So, you know, currently hardware is identified as a critical component, so when you look at sensors for examples, from a software perspective it's garbage in, garbage out, you know, if you have bad sensor input. So what does the panel think about leveraging some of the existing safety standards to sort of be a stopgap with dealing with hardware from a purely cybersecurity point of view until we mature to that point?

DR. VASSERMAN: Can I answer this?

MS. RICCI: Sure.

DR. VASSERMAN: Since mine is not a question, it's a comment, there was a particular processor that had a very interesting bug in it where if you made a specific system call with a specific set of options it would kick you out to root whatever process you had. Ask me later; it's really cool. But if you have no way of knowing which device ended up with that processor, this isn't even something you can fix in an operating system, you can mitigate it but you can't fully fix it without replacing the chip.

MS. RICCI: Totally understand. So our time is up, so unfortunately, we're not going to get a lot of answers to the last question, but great points. And we have time for a 30-second answer.

MR. CORMAN: Dr. Julian Goldman is in the room here. I love his term when he does his workshops, it's let's focus on preventable harm, and with the very narrow initial scope that a lot of the folks up here are driving, there's a lot of preventable harm we can work on and then we'll work on the harder stuff.

MS. RICCI: So I'd like to thank everyone on the panel. I think it's been a very informative session. Thank you.

(Applause.)

DR. D'AMICO: Okay, everyone, we're going to move to the last activity of the day, which is the breakout panels, breakout sessions. So as a reminder, Aftin said that we're going to combine the breakout session from earlier today with the one now due to the early dismissal by OPM. So if everyone would do me a favor and please look down at your name tags, you will see that you've been assigned a certain breakout session number and that breakout session number also coincides with the location of your certain breakout group.

So for a little bit of orientation, Group 1 is going to be over here, and Group 18 is going to be over here. If you have Group 19 as your breakout group, that is going to be in Room 1506, which is behind us, and if you are Group 20, 22, or 21, that is going to be in Rooms 1404, 1406, and 1408, which is behind all of you over here. So we're going to take a quick break, and we plan to reconvene at 3:20.

DR. ROSS: And I have the breakout question folders for Mari, Ross, Sega (ph.) and Joseph Cody, so if you haven't picked them up, please come up. The way we are going to divide the questions is that Groups 1 through 10 will do what was going to be the first breakout, which is everything but CBOM, and Groups 11 through 22 will do the CBOM-related topics.

(Breakout session from 2:15 p.m. to 3:14 p.m.)

(Whereupon, at 3:16 p.m., the meeting was continued, to resume the following day, Wednesday, January 30, 2019, at 8:30 a.m.)

CERTIFICATE

This is to certify that the attached proceedings in the matter of:

PUBLIC WORKSHOP - CONTENT OF PREMARKET SUBMISSIONS FOR MANAGEMENT OF CYBERSECURITY IN MEDICAL DEVICES

January 29, 2019

Silver Spring, Maryland

were held as herein appears, and that this is the original transcription thereof for the files of the Food and Drug Administration, Center for Devices and Radiological Health.

____

TOM BOWMAN

Official Reporter