INFORMATION SECURITY AND PRIVACY ADVISORY BOARD

Established by the Computer Security Act of 1987
[Amended by the Federal Information Security Modernization Act of 2014]

MEETING MINUTES

March 15 and 16, 2018

American Institute of Architects
1735 New York Ave., NW, Washington, DC, 20006

Board Members	**Board Secretariat and NIST Staff**
Chris Boyer, AT&T, Chair, ISPAB John Centafont, NSA Laura Delaney, DHS Greg Garcia, Health Care Coordinating Council Patricia Hatter, Barrick Gold Corporation Marc Groman, Privacy Consulting Brett Baker, Nuclear Regulatory Commission, OIG Jeffrey Greene **Absent with Regrets**	Matt Scholl, NIST, DFO Robin Drake, Exeter Government Services, LLC Warren Salisbury, Exeter Government Services, LLC

Thursday, March 15, 2018

Mr. Matthew Scholl, Board Secretariat, opened the meeting at 9:06 a.m., Eastern Time in the place of Mr. Chris Boyer, Chair who was called away.

On behalf of the Board, Mr. Scholl welcomed Mr. Brett Baker as its newest member. He comes to the Board from the Nuclear Regulatory Commission (NRC), Office of the Inspector General. He is not just a member of the agency but also from the inspector general community.

Welcome and Remarks

Dr. Charles Romine, Director of Information Technology Lab, NIST

Mr. Scholl welcomed Dr. Charles Romine to the meeting to brief the Board on activities happening with the National Institute of Technology (NIST) Information Technology Lab (ITL). Dr. Romine thanked the Board for its service and emphasized the importance of its work.

Most of the type of security activities and privacy activities the lab undertakes are much larger than cybersecurity. Of the lab's seven divisions, two are dedicated to cybersecurity and privacy. Four of the other five divisions engage in cybersecurity activities. During the last year, the ITL undertook an exercise in developing an understanding of its role in what NIST does. Its purpose for existing is cultivating trust in IT and metrology.

Information metrology is a critical component in mathematics, mathematical modeling, and statistical analysis to improve the state of metrology in the rest of NIST. Dr. Romine

provided an update on key NIST programs. Two foundational areas of activity are risk management and cryptography. The lab started in 1972 and the lab's 50th anniversary is coming up. NIST pioneered risk management in the very early 2000's, before the FISMA Act of 2002. There is also a long history in identity management and privacy engineering. Within the last five to ten years, there has been an emphasis on trying to make privacy less complex and trying to understand how to provide tools and the ideas associated with privacy issues. The other major areas are the National Cybersecurity Center of Excellence (NCCoE), and the National Initiative for Cybersecurity Education (NICE).

Mandates have also been important for ITL over the years. Mandates may come from Congress, or the administration. Various administrations have assigned NIST important goals in the areas of cybersecurity. A range of legislation governs what NIST does as well as executive orders in which NIST responds.

The most recent legislation is an update of the original Federal Information System Modernization Act (FISMA), and the Cybersecurity Enhancement Act of 2014. A recent executive order involves strengthening cybersecurity federal networks and critical infrastructure. NIST is heavily involved in activities under that executive order. The legislation that included establishment of the NCCoE also included an update that combined the NCCoE with the National Strategy for Trusted Identities in Cyberspace (NSTIC). Those programs merged two years ago.

NIST has a role in the Help America Vote Act (HAVA), with the security of voting systems as well as their interoperability, usability, and The Computer Research and Development Act (CRDA) of 2002. Those were some of the legislative activities and executive orders that govern what NIST does. More mandates are expected in the near term. While being considered able to handle the tasks being requested, there is a capacity issue. There's a finite set of resources and when people are already working many hours a day to meet some of these additional mandates, taking on additional work with current resources can be a challenge.

Additional money is rarely available with these mandates or orders and most come from committees on the authorization side to request or order what will be done. NIST does request funds to support additional work, but often it doesn't happen and NIST must perform the tasks anyway. NIST must prioritize what un-mandated things are most important, and what can be turned down for a period of time in order to answer what's needed immediately.

There's a request for NIST to support small business more actively. There's some concern about the evolution of the Internet of things (IoT) and its implications. IoT security is a major issue, along with workforce development in the area of cybersecurity.

Other updates include the Cybersecurity Framework (CSF). The framework for improving cybersecurity critical infrastructure is already one of the most successful work products that NIST has created relating to cybersecurity. It is an answer to Executive Order (EO) 13636 from 2013. Version 1.0 was released exactly on schedule despite a short-term shutdown just prior to the release.

It was then codified through the Security Enhancement Act of 2014. The second round of the framework, version 1.1, was released in 2017 and the final release of version 1.1 will

happen in spring, 2018. The NIST Cybersecurity for IoT program has a workshop scheduled at RSA in April, 2018. Dr. Romine will be there along with many other people from NIST.

The way to best cultivate trust is to address definitions, guidance, and best practices associated with IoT security and privacy, reference data and some software tools to be developed in the space, and coordinate standards for the digital economy. NIST released the "Interagency Report on Status of International Cybersecurity Standardization for the IoT" in February.

One of the tenets of Executive Order 13800, called upon the Department of Commerce (Commerce), in partnership with the Department of Homeland Security (DHS), to address the actions that can be taken against botnets and other automated threats. That report was released by Commerce and DHS in draft and comments sought.

Work has grown substantially in post quantum cryptography, both because of internal input from NIST, and appropriation from Congress. Quantum computing has the potential to make the current public key cryptography systems that secure financial transactions completely irrelevant. NIST made a formal request for algorithms for post-quantum cryptography from the cryptography community.

Once the period for submissions closed, sixty-nine solutions were received. The group is now in the process of analyzing them and seeking input from the community to determine whether they are first resistant to classical attacks; it doesn't do any good if it's quantum resistant but not resistant to classical attacks. There was a mixture of submissions from very large teams and a few individual contributions. Some of the entries have been determined to be not quite adequate. These competitions are completely open and transparent. The proposers are visible, algorithms are visible, and the attacks that are used to determine whether something is acceptable or not are visible. Participants came from twenty-six countries, representing six of seven continents.

The NICE initiative has turned out to be one of the most important foundations of cybersecurity workforce that is being undertaken by the federal government. The location for the NICE conference coming this November will be announced. There is also a subsequent education conference in the K-12 arena to be held in San Antonio, TX. The first annual cybersecurity career awareness week was held last November.

NIST Manufacturing Extension Partnership (MEP) and Cybersecurity

David Stieren, NIST

The Chair welcomed Mr. David Stieren to the meeting to brief the Board on the NIST Manufacturing Extension Partnership (MEP). Mr. Stieren spent his whole career at NIST supporting manufacturing through metrology, through standards and most recently for the last ten years at NIST MEP while working with its nationwide network of centers to provide manufacturers, especially small manufacturers with the hands-on assistance that they need.

MEP is not a research and technology development organization, but a national program that's a public-private-state partnership managed by NIST. MEP provides hands-on technical assistance to the nation's manufacturers. Its legislation allows it to work with manufacturers of all sizes. The requirement is they have operations in the United States. Many of the companies served are also global.

The majority of the manufacturers that MEP serves are small manufacturers. Of 347,000 entities that exist in the U.S. that are considered manufacturing organizations, about 99 percent of them, are considered small to medium size enterprises according to the Small Business Administration definition. This means there are 500 or fewer employees at a location. About eighty (80) percent of companies fall in the 25 - 250 person range. When MEP was created in 1988, most manufactured items came from Japan. Those Japanese products were cheaper than the U.S.-made products they competed against and they were higher quality.

An economic study determined that the demographics of U.S. manufacturing sector included a preponderance of small manufacturing companies. These companies were at a competitive disadvantage to their foreign counterparts because they didn't have access to the information to stay up to date with the most current technologies, processes and approaches. That's why MEP was created.

Multiple economic studies have occurred since 1988, and they indicate a market failure remains. Small manufacturers are very good at what they do. They are very agile, but focus on meeting payroll, and the next contract. MEP goes into companies and acts as trusted advisors. The advisors help those companies understand what they're doing, where they need to get to, and then help them get there. Sometimes it's through hands on services or through services MEP connects them with. They can also be referred to partners either at the state or local level or consultants MEP helps them work with, through its national network.

There are centers in every U.S. state and Puerto Rico. MEP has 1,300 people who act as trusted advisors. These are people who have deep experience in manufacturing. They go in and work with companies. MEP partners with about 2,100 service providers and third parties every year. MEP centers are either parts of state government agencies, or they are embedded within universities, typically an extension component of a university, or they are non-profits. MEP center staff are non-federal staff. The only federal component of the MEP national program is the national program office at NIST.

The 2016 operational budget was $130 million. Companies work with MEP centers. MEP is not a grant program. MEP people interact with many companies every year. When a company interacts with an MEP center, the company pays for the services. It's a very high impact business model for companies. When a company invests in the process, they'll make the change they seek. When companies receive grants, the changes tend not to stick. In return for the $130 million federal contribution, at least a dollar of non-federal money must also go into it. Companies pay and states often provide money as well toward the fees.

MEP advisors work with companies one at a time to go in and provide them hands on assistance. Centers are well connected with the manufacturing base in the states. Annually, MEP works with around 30,000 U.S. manufacturers, and conducts about ten thousand detailed projects with manufacturers on a national scale.

There are about 45 people in the program office at NIST. There are five regional managers stationed around the U.S. who are NIST employees. They handle day-to-day operations and management on behalf of NIST. It is a national network but a NIST program. Centers may have different areas of expertise, but assist wherever needed. Supply chains are an

enormous priority. When advisors talk to manufacturers about cybersecurity, if they know about cybersecurity, it's because there is a supply chain requirement.

They provide service in these areas: process optimization, Baldridge designation preparation, technical manufacturing services, and product/process development, innovation, supply chain development, work force development, marketing, IP management, financing/access to capital, sustainability, export, and market diversification.

They work with companies in several stages beginning with initial engagement, which happens in a number of ways. Those who are interested in working with an MEP center begin with an assessment. Companies provide information on their current state. MEP will identify gaps and determine the company goals. MEP finds twenty-five to two hundred fifty employee companies are best to work with. They work with startup and very small companies as well. MEP goes into manufactures and provides hands on technical assistance.

They served twenty six thousand companies in 2017. MEP assesses impacts at a follow-up interview six months later. To date, impacts listed from 2017 include $12.6 billion invested, $1.7 billion in cost savings, and 100,000 new or retained jobs. Eighty percent of companies are established manufacturers. The number of new or startups is very small, less than ten percent. MEP has good working relationships with SBA, and portions of the Department of Defense (DoD).

They have been partnering with the University of Delaware via the SBA for cybersecurity. They developed a good basic assessment tool for cybersecurity that is available to the national network. Follow-up interviewers ask why companies hire MEP. MEP is a public good national consultancy. Their Net promoter score is 83%. It rates the likelihood of companies recommending MEP to other companies. MEP averages 80+ percent annually.

When asked, participants cite many business challenges, but cybersecurity is not named among them. Cybersecurity is not part of the small business culture. MEP is trying to help small manufacturers implement cybersecurity, as in the framework, to protect themselves. Great things have been brought forth from NIST in terms of the framework. Businesses tend to view new government requirements as eating into profit. But when cybersecurity is viewed in terms of asset protection, the perception of cybersecurity changes.

There is a small group of centers working on developing a cybersecurity practice to assist others. About half of the centers provide cybersecurity assistance. The goal is to help companies understand cybersecurity requirements, what they mean, and how to implement them with the framework. It means taking guidance and reducing it to practice for small manufacturers.

De-identification and Anonymization

Jason Suagee, NIST

The Board Secretariat welcomed Mr. Jason Suagee of NIST to the meeting to brief the Board on de-identification and anonymization. The work being done at NIST resides within a different division of the ITL. The presentation provides an idea of what de-identification is, what it does, and how it is done. The objective is to understand the work that's taking place so that as new standards come out, conversations start more and more about de-identification, anonymization, and synonymization, what work is taking place to help that

discussion and help that technology? Classical de-identification techniques work, but are not effective if one has a whole lot of computing power and can just reengineer the database.

Differential privacy can be used as an implementation for a more secure de-identification technique. Difficulties exist working with differential privacy or any other kind of similar privacy ideas, because answers to queries have noise added to them. Random numbers generated in some range are intended to obfuscate the original data down to the level of an individual record. The whole idea is to protect against identifying who is in a dataset down to the individual, then privacy is protected. When noise is introduced, the data mostly becomes unusable. Over the past year and a half to two years, NIST has been looking into various ways to try to implement existing algorithms or develop new ways of doing de-identification using differential privacy methods.

With de-identification, it is potentially possible to re-reverse engineer ridesharing data to calculate possible members of the dataset down to the individual person. Famous celebrities have been identified in this manner and it was not that hard to do. De-identifying these datasets using differential privacy methods is difficult, as data could be in strange formats. It was mainly about replacing identifiers with pseudonyms or hashes of the real data and this type of de-identifying does not work well. From a legal point of view, there is a point where de-identifying data is sufficient. This standpoint is problematic as the legal requirements are not strong enough and there is no way to determine what "strong enough" would be.

It could be done by hiring some white hat hackers and try to reverse engineer who is in the dataset. That's been done a couple of times with certain datasets, but reverse engineering is not a foolproof method. In different privacy paradigms, there are various parameters, like epsilon and delta, which quantify the privacy lost or the privacy budget. It is not known exactly what those mean in real world terms. What privacy law corresponds to what value of epsilon?

Apple was playing around with differential privacy recently when they set an absolute epsilon value of one for every day, for every person. Most researchers who use differential privacy try to limit the epsilon to less than one for every dataset. The dataset could be huge so, no one knows quantifying privacy loss is difficult. It is the probability of being re-identified. The focus should be more on ways of producing differentially private datasets, de-identified datasets, and specifying one of these parameters as input and not on what it actually meant.

A lot of privacy research has progressed to models that are more mathematical with a lot of probability and statistics. There is a large disconnect between the research and its actual application. A lot of researchers focus on one-dimensional datasets, datasets with a single data type, and most data in the world has more than one type. It becomes very complex.

What's really needed is using differential privacy instruments to state an epsilon limit that can't be exceeded. It becomes a quantified budget to stay within, then someone runs software to de-identify a dataset to get a reliable result within a set epsilon. It may not be theoretically possible for most datasets. A theoretical barrier might exist where the dataset can't be released as it can be de-identified, but it may be useless. Many of these theoretical barriers make it seem improbable these methods will work.

There is the complication that datasets usually have many fields that are complex and most research deals with real value data or simple datasets. For instance, even the Netflix dataset, has approximately two thousand attributes. The attributes show if a person watched a movie or not. This is actually a simple dataset since it's binary, but this is still hard to de-identify. There are other paradigms or generalizations of differential privacy, like probabilistic differential privacy, say that these are relaxations of the differential privacy paradigm, which could allow for more utility in the end dataset.

Individual differential privacy is interesting as it could result in a much higher utility, but at a small cost in terms of what is being shown when a de-identify dataset is de-identified that cannot be re-identified. It is hard to define the trade-off, but it comes at a cost in terms of what differential privacy, might legally claim individuals in the data set cannot be identified. Legally, it is not certain that an individual is in the dataset, but a person could infer certain people are in the dataset.

The more data is de-identified, the less potential benefit of the dataset. The probability of re-identification needs to be quantified. Every agency in the federal government is asked to do this and uses different programs so the answer may not be organized at this time. The goal is to get to a level where datasets are uniform and create a software package, which takes those algorithms, and maps to the different kinds of datasets, and offers results.

Mapping algorithms to different kinds of datasets is a hard problem and will likely be around for maybe 15 years at least. There probably is not a solution, therefore people may have to sacrifice a lot of their privacy for research purposes.

There's a legal approach across the government, so it's being done in a formal way from agency to agency, knowing full well that there's always risk. There are the HIPAA regulations, which provides guidance, but there is not any regulation that provides good guidance on differential privacy outside of advanced mathematical computer science. Some uniform technique must be applied that these agencies use to affect their datasets. It might be a reasonable goal so that in the future there is a uniform software product, which does a certain job and is streamlined in an algorithmic way.

There are only a few software companies that claim to handle differential privacy, one of those, Immuta, which is being tested at Lincoln Labs at MIT, and they expect to do a test run with their software. The software may be the interim "holy grail" of fixing the government's problems. Then there must be a team of people who can adapt the software to the datasets that the government has to de-identify. It is a level of expertise that is needed.

The other issue is releasing synthetic datasets, a lot of privacy research does not focus on this. It focuses on research. When releasing synthetic datasets, which have the same properties as the original datasets, it becomes much harder. Maybe government agencies could switch to that paradigm for some datasets, such as health records from NIH, access could be granted to the dataset through this portal. Receiving answers in this way could actually protect a quantifiable level of privacy.

Government Cybersecurity Assessment and Risk Tool (GOVCAR)

Patrick Arvidson, Department of Defense (DoD)

The Chair welcomed Patrick Arvidson, from (DoD) to the meeting to brief the Board on the government Cybersecurity Assessment and Risk Tool. It is a mechanism being used and

explored across the government to understand, quantify, and have a better understanding of where risk decision should be applied.

There is a trend where dependencies are increasing, missions increasing, and networks are getting larger. The cyber threat level continues to increase as well. While all of these issues continue to increase there is a decline in cybersecurity resources to combat threats. There is a tendency to treat these problems as a one-to-one problem, meaning there is one thing that will resolve all issues. This is rarely ever the case. When it's treated as a one-to-one as opposed to a many-to-many to many problem, a solution based on a single perspective may not be the solution for all concerned. It is more of a tiered approach; meaning, a holistic viewpoint is needed.

The first step is to explain how to talk about the threat. There can be conversation between architects, engineers, operators, business analysts, and the system administrators, researching through years' worth of incident reports to try to verify what they were trying to do in each report. It means checking validity of past reports and investigating how long attacks have been happening, and determining how to protect or what the response should be.

Those capabilities were put back in the framework, against all the things that the adversary does, then the gaps were analyzed. The data was analyzed to find what was lacking and where there was overinvestment.

Considerable effort was expended to know how to do this kind of a risk assessment as threat based versus compliance based. Executive Order 13587, to assess the national security system, is moving to this type of a system looking at it from a cyber-operation awareness perspective. They decide the locations to work with, show them how to build a framework, build that framework. They assess where the risks are. The team comes in and does an inspection. The team goes in and hacks the system as if it is an agent, based upon a heat map they have.

A big win came from a small base out of the Air Force, which was an external learning system with an external website that is externally hosted. They do not have many resources, but they found the information and followed the process. They are now getting an award for small base cybersecurity. They have had no incidents since they started a year ago. They do not have any systems put in on the base, they just follow a base process.

The team is very open and shares what it produces in order to be as transparent as possible. They take the threats and provide it in a useable format or provide a process to counter as a threat. GOVCAR provides the assessment information. It provides the glue that makes the CSF and the department go.

They use the "big system" that everyone uses from the Defense Security Service (DSS), National Security Administration (NSA), and DoD and rebuild the heat map for them, as they do not build until they need it. They use the same proprietary system with a few tools. They are building the same process for the perimeter. There is a family of tools coming out, the first one simple, and points out how the adversaries can get into the system and what to do to prevent that. It is called NexGen tools officially, it is a tool that preloads a heat map for, and shows where, all the adversaries are and gives a list of capabilities. Most of the

capability comes from the network. Generally, when there is a list of capabilities, the list of capabilities are tied to Special Publication (SP) 800–53.

The next tool is more in-depth and designed for engineers to go through and understand the trade space. It is an analytical network and the tool says if there is the capability to do tactic analytics design within the network, to build that to protect and respond to analytics as well. The group is in partnership with another company that is investing its own hours to build a security situational awareness tool that will provide real-time feeds, based on everything that is coming in from internal equipment.

GOVCAR is also partnered with foreign partners, North Atlantic Treaty Organization (NATO). In places the tools have been used, there has been significant improvement in security and significant understanding of security from people whose job function is not security. The mindset is to reorient the entire system security engineering process towards this model. Tools will be developed, shared with the whole government and everybody. DoD has provided funding. DHS will also receive information.

Artificial Intelligence (AI)/Machine Learning Bias

Karthikeyan Natesan Ramamurthy, IBM Research

The Chair welcomed Mr. Karthikeyan Natesan Ramamurthy from IBM Research to brief the Board on artificial intelligence and machine learning bias. IBM Research has more than two dozen researchers working on different problems from AI, quantum, industry, and physical sciences, Mr. Ramamurthy presented a broad picture of a network with its components such as sensors and people trying to accomplish a task.

Generally, an AI and machine learning flow would look something like this process: sensors send data to model, the model sends predictions to actuators, and actuators send actions to the world. Sensors could record things like various attributes, such as the performance, and various characteristics about the users. What sensors receive as input provides the data on people and tasks. The data becomes a management database in effect. Models can be developed in order to accomplish specific tasks. Any AI application can be thought of this way and fit it in this model.

IBM Research has a science for social good initiative. Every summer, for the past three summers, IBM has been calling students and asking them to work on projects that provide social benefit for people. Some of the projects that IBM has done involved how to accelerate scientific discovery, how to diagnose cognitive disease, and how to find inspiration in nature. People talk quite a bit about artificial general intelligence (AGI), but it's a far way off. The kind of AI one sees in science fiction films is not going to happen any time in the near future. They look at the world as it is now and see what problems can be solved now.

Things like mission level safety, security, bias, and privacy are real and happen now. Technology for the sake of humanity could provide amazing benefits and important impacts. It is why the whole discussion of bias and ethics is super important. Mr. Ramamurthy provided some examples of where bias happens in AI. These examples are based on well-defined studies that have been published on bias.

"Discrimination in Online Ad Delivery", a paper published by LaTanya Sweeney of Harvard University, provides an example of algorithm bias. The paper investigates the delivery of types of ads based on the type of name entered. As described in the abstract, she used a

sample of racially associated names, and found statistically significant discrimination in ad delivery on more than one website. It was found that names given more predominantly to black babies were much more likely to generate ads suggestive of an individual having an arrest record even if the individual had no criminal record at all. Names determined to be more Caucasian generated ads that were neutral, and did not imply an arrest record even if an individual had an arrest record. Other examples of machine learning include a chat bot that Microsoft developed and launched in Twitter. Very quickly, it picked up very offensive language, and Microsoft had to shut it down in 16 hours.

An application which predicted recidivism was used in a few counties in the United States for figuring out whether a defendant will re-offend. It consistently provided high risk scores to African Americans compared to Caucasians. The actual truth is that the predictions did not match up with what happened afterwards. The Caucasians who got the lowest scores were most likely to commit crimes. Whereas the African Americans who got highest scores did not reoffend. There was a very evident bias there. The other very interesting example is predictive policy. Empirical evidence tells us that police officers, whether implicitly or explicitly, have some sort of racial bias. The conclusion is based on algorithms used to decide which neighborhoods to patrol. The algorithms tend to reinforce bias because if the algorithm shows a particular area has more crime, it will send more police officers. If more police officers patrol an area, they will find more crime. It is a never ending circle. The algorithm does not correct for bias or make allowances for the actual statistics for an area.

The next example of bias is in Word Evidence. As an example of what word evidence is, if one wants to do machine learning tasks in a natural language, words need to be converted to numerical representations and then sentences. Ways exist to convert words to numbers. These are statistically learned using a large base of historical data. If one takes the numerical representation of the word "man" and subtracts it from the numerical representation of the word "woman", it's equal to subtracting infinitely. The result is to portray woman as a lesser value than the value for man, or man is to woman as king is to queen. It can also give results such as, man is to woman, as computer programmer is to homemaker.

The papers talk about how these results come about. Examining the words "he" and "she", mostly the "she" occupations, it lists all of the stereotypical things about what women can do. For the "he", words like mason, skipper, pro DJ, philosopher, others. It was proved it happens not just for adjectives but also verbs, occupations, etc. The essential problem is proper training data was not supplied to train the algorithms. Inequality is one of the reasons why people should care about bias in AI. Giving the right values to AI will create trust. It's important to involve people to learn about proper solutions that satisfy everyone. Sensors have limited capability to sense the world. It leads to biased data. Actuators, or people who take actions based on the model, can have a limited mandate. It was shown that, when people from African American communities try to rent Airbnb, the approval rates are lower. They created an instant approval, in Airbnb. It doesn't have to be approved by anybody.

There is no one definition of bias. There is proxy bias, presentation bias, statistical bias, and inductive bias. All models are approximate representations of the world. There are possibly 100 different manifestations of bias. There is really no one definition of bias to rely on

because it manifests in many different ways. Unacceptable bias in the data must be removed. People have come up with more practical definitions. It takes note of situations when the bias does not throw in equal benefits to all groups of a population, meaning the benefits for one group of a population are substantially higher or lower compared to another group. The same sort of principle applies with fairness. The optimal bias exists all along, but this leads to the principle of fairness because it says when unacceptable bias is removed, there are fair models and fair actions. Bias is a concept. Fairness is a goal. Protected characteristics can be used to benefit others as in doctors making diagnoses while considering race, because certain races are more likely to have certain diseases. There are variances of bias and fairness. This is beneficial to everyone.

There is a fine line between regulation and being fair to everyone. There are two types of fairness. Outcome and process. Bias can enter at different points. Processing the data before it gets into the model, and removing bias from the data are processing techniques. And if one wants to process the model, change the model to make it less biased, it is in processing. Processing to remove bias in the predictions, is called post processing.

Trusted Internet Connections and High Value Asset (HVA) Program Update

Crystal Jackson, DHS

Sean Connolly, DHS

The Chair welcomed Ms. Crystal Jackson and Mr. Sean Connolly, both of DHS, to the meeting to update the Board on trusted internet connections and the High Value Asset program. There were four different stakeholders at the initiation of the trusted internet connection (TIC) program: the Office of Management and Budget (OMB), which sets the policy and strategy for TIC; DHS, responsible for the security baseline and compliance measurements; Managed Trusted Internet Portal (MTIP), the program where agencies that don't support their own access points and their own security stacks; and the agencies. Executive Order 13800 tasked the Director of American Technology Counsel to coordinate and report to the President, OMB, DHS, and the General Services Administration (GSA) on modernization of federal IT including improving the security posture of agencies. The report offered recommendations towards two categories of efforts. The current policy requires two access points for redundancy and if the agency wanted more access points to make a request through us or OMB and give justification. One of the goals is to rapidly increase the number with greater than two trusted connections, and determine what that means for agencies that are supporting cloud versus virtual cloud.

DHS, TIC PMO (Program Management Office) will host a series of working groups for agencies to discuss the architecture that has been envisioned. The team is taking lessons learned from the pilots and use those both with the TIC program office and national security deployment (NSD).

The HVA program was established back in FY16 as a result of breaches that occurred throughout the federal government. DHS and OMB determined that agencies need to identify what their most critical IT assets are. DHS was then tasked to provide a series of high value asset assessments to determine how those assets may affect the ability to either serve the public, the United States or the federal government. A series of OMB memos were generated directing certain actions regarding agency TIC. A directive was provided for DHS

to help agencies identify high value assets, but to orchestrate, coordinate, and conduct high value asset assessment against them.

The job in doing the assessments is really to identify the impact a high value asset offers to the federal government or to the American public. A report is generated, provided to the agencies with recommendations for the agency to mitigate the risks.

For FY 18, DHS is working to do 30 risk and vulnerability assessments across the federal government and 30 security architecture reviews. They are on target this year to make the goal. At this date, 10 security architecture reviews are done, with another three scheduled to finish before the end of the month. Nine risk and vulnerability assessments are completed to date, with another two scheduled to finish before the end of the month. The risk and vulnerability assessors come in and do the remote test one week, followed by an onsite test the next week, so the tests are done within two weeks. The security architecture review makes an evaluation from a business risk perspective.

Having a community of interest really brings everybody in to share challenges and to discuss those as a group to be aware of dependencies. There are many Chief Information Officers (CIO) and Chief Information Security Officer (CISO) councils, where things are being shared. The hope is all this information will be directly applied to the HVA program. The goal is to secure high value assets then also secure the remaining systems. There is a draft strategy and process where it outlines what the next three to five years will look like from an HVA standpoint.

The approach of identifying the most critical information systems, developing visibility in to the cyber security posture of those systems, and ensuring that whole government approach to effectively secure and manage the risk of those systems is key. As this approach continues, things will be in a much better position than three years ago, to be able to identify whether there is a major vulnerability out there that has a negative impact to the rest of the agencies.

The team has worked with agencies that have self-declared HVAs, but upon assessment noticed that the asset was not an HVA. With the architecture review and the vulnerability assessment, it is not just HVA specific but also the enterprise level. There is conversation with the agency because the agencies know a lot more about their systems and they provide rationalization on why it was a high value asset.

GAO Report 18-211 CRITICAL INFRASTRUCTURE PROTECTION: Additional Actions Are Essential for Assessing Cybersecurity Framework Adoption

Nick Marinos, General Accountability Office

Mike Gilmore, General Accountability Office

Kush Malhotra, General Accountability Office

The Chair welcomed Mr. Nick Marinos, Mr. Mike Gilmore, and Mr. Kush Malhotra, all from GAO, to the meeting to update the Board on GAO Report 18-211 on critical infrastructure protection. The Information Technology team is comprised of four directors. The group represented here is the Cybersecurity and Information Management Issues Team. The team looks primarily at critical infrastructure as it relates to cybersecurity issues. They also deal

with privacy related issues or issues where the federal government may be using or sharing sensitive data

The Cybersecurity Enhancement Act tasked GAO with the creation of bi-annual reports. The first report looked at promotion and development of the Cybersecurity Enhancement Act. The group is currently working on its second report, on adoption of the framework. The final two reports are intended to look at the success of the framework in better protecting critical infrastructure from cyber threats.

One of the challenges that was encountered early in this report, is there is no true definition of what it means to adopt the framework. The framework is intended to be flexible and adaptable. Flexibility is a great idea for creating something that could be used across all critical infrastructure sectors, but it presents a challenge in trying to measure how it is used.

Of the sixteen critical infrastructure sectors, twelve have been identified as taking steps to develop guidance for their sectors to adopt or implement the framework. Of the twelve, the energy sector, the financial services sector, and the healthcare and public health sectors went further and came up with more specific guidance or tools. The energy sector developed the Cybersecurity Capability Maturity Module (C2M2), a detailed tool for use within that sector to determine cybersecurity preparedness. It demonstrates the implementation of different aspects of the framework.

The financial services sector decided to create and add a sixth function. They also developed the draft cybersecurity profile, and a mapping of the CSF to various regulations within the financial sector. The healthcare sector, primarily through ITrust, mapped all the various security controls and the privacy controls that healthcare entities have to follow to the cybersecurity framework.

The challenges to using the document framework were entities had limited resources to commit to implementing the framework along with the proper skillset and expertise. They did not have the right people or the most knowledgeable people. Another issue that was noted across a few sectors had to do with regulations or existing requirements, which predated the framework. There was also the issue with how relevant cybersecurity is to everyone across the board. While it is important to all of us, it is not always equally important across the entire landscape. Cybersecurity was just not something that was a focus in terms of how to commit sector resources and efforts.

The team hopes to get Board input on protecting critical infrastructure from cybersecurity threats. It is a tall task for GAO and it makes it somewhat similar to the views of how best to evaluate whether what is being provided, assistance-wise, from the federal government to specific sectors is the right formula for that assistance.

The evaluation looked at version one and the framework. Version two is moving away from the concept of tiers because during the public comment period, there was a perception that this was going to be a regulatory evaluation tool. One of the challenges has been when GAO or government agencies ask a question, it is a self-fulfilling thing because a lot of companies track really closely what is going on in Washington who will inherently respond to those questions because they know what is going on.

Review of Thursday Briefings

The Chair noted AI and machine learning to improve cyber security is going in the right direction, but turning over the operation of critical infrastructure to a reasoning engine that is learning could be sabotaged resulting in a less secure environment. The Board needs to think about how that will be dealt with, how to identify it, and how to alleviate the issue. This brings along a significant privacy issue and has implications on how the government will use the data.

The issues that concern the Board are not new issues. What is new is the misperception and misunderstanding that AI and machine learning produces factual and accurate results without bias. Bias can never be fully eliminated, even though that is the goal, but making sure policy makers know that while the data being produced from algorithms are not necessarily wrong, it should not automatically be considered right. As these algorithms and machine automated learning continue to be tuned, some entity needs to be at the other end and confirm whether there is suspected malicious activity or whether it is in fact malicious activity.

Public Comments

No public comments were presented.

Meeting Recessed

The meeting recessed at 4:17 p.m., Eastern Time.

Friday, March 16, 2018

The Chair opened the meeting at 9:03 a.m., Eastern Time.

Internet of Things Cybersecurity Standards Report

Lisa Carnahan, NIST

Michael Hogan, NIST

The Chair welcomed Ms. Lisa Carnahan and Mr. Michael Hogan of NIST to the meeting to brief the Board on the Internet of Things Cybersecurity Standards Report. NIST, with the Department of Commerce, undertook the effort to form the task group. Ms. Carnahan convenes the group, but works in the NIST Standards Coordination Office.

Mr. Hogan, the NIST ITL Standards liaison, served as the primary editor on the document for the Interagency Work Group on International Cybersecurity Standardization. It is an interagency work group that focuses on international cybersecurity standardization. It was formed out of a specific recommendation out of NIST Internal Report (IR) 8074, which was the document with recommendations on how to coordinate international cybersecurity standards in the federal space.

The work group formed in October, 2016. There are approximately 78 members representing 29 agencies. The group tries to meet three times a year. It does topical discussions on a wide range of things. People also attend from Standards Development Organizations (SDO) and private sector representatives also provide industry perspectives.

The document was formed under the IoT task group. It was well received and had already been publicly reviewed. It speaks to looking at application or sector areas and then looking at standards from a core technology area perspective. It also looks at individual standards and their maturity level in the development standard process and the level of adoption. It was very well received and approved.

The working group approved Kat Megas from ITL as co-convener with Mike Ruza from DHS. It was a 90-day effort. There were 54 people on the mailing list, representing about 13 agencies. The task group met from June to November. The decision was made to extend the working period from the original ninety days. The team wrapped up in November with a document. The working group met in December. The group agreed to release the document for public comment and publish as a NISTIR. It became NISTIR 8200. It mirrors NIST IR 8074 for IoT, cybersecurity Standards in IoT. The public comment period lasts until April 18th.

The group wanted to make sure it didn't duplicate existing reports, but there was no one focused on the state of international cybersecurity standards in IoT. Section 4 talks about what's an IoT component, what's an IoT system, what's an IoT environment. There are use cases or application areas. There is a nice spectrum: Connected Vehicles, Consumer IoT, Health and Medical Devices, Smart Buildings, and Smart Manufacturing.

In 2015, cybersecurity was divided into ten areas, and an eleventh, hardware security was added later. The group looked at the cybersecurity objectives for IoT in general and for

those areas. The group also looked at risks and threats for those areas. The group felt it could remove industrial control systems. Much work has been done, and added IoT systems for many areas to its list of threats and risks. It was a matter of extending existing work to the concept of IoT in general.

The standards were annotated in Annex D. The report is approximately 180 pages. There are 103 pages of the Annex. The report itself is about 60 pages including the conclusions.

The task group worked on looking where might there be gaps that would warrant a new standards initiative. "Unknown" appears in the report in several places to define market impact. It's not as clear-cut in these instances that standards have been taken up by the market place.

The definition of "uptake" means at least two or more implementers are pushing product in the marketplace. The group was very eager to get industry and SDO input. Follow-on activities started last year following handing off the report. The working group focus was determining if the team wanted to open the report for public comments. The outcome was to get public comments. It was decided to use NIST's interagency report publication process. There was a final call for the working group to make comments on what the task group had done, and then it went out. Comments are due in roughly 60 days, on April 18th. The document was greatly improved by input from the private sector. There have been three sets of comments to date.

Proprietary and open source were taken out, because the guidance for the work is the OMB Circular A-119, Voluntary Consensus Standards. In reality, the market that's being offered and consumed, is a mixture of proprietary, voluntary consensus standards, and open source. That's what people need to understand overall, but a key aspect is the portfolio of Voluntary Consensus Standards.

The agencies would be interested in doing some project proposals that would be specific somewhere. However, it's understood that patches can't be pushed out to a lot of IoT components, as they're not set up for that. For the radio frequency identification (RFID) technology, they've done a set of standards that would deal with what might be found in an RFID deployment. It's not clear when some of these things might pick up steam and start being used. Market uptake for more than advanced encryption standard isn't known.

All the network security standards need to be revisited, because the way some IoT systems are set up. Ownership of components can get fuzzy. It's not like a traditional IT system doing a risk assessment. The group should look at Bluetooth, transport layer, and security because of the attributes of some IoT systems. The cryptographic technique world of standards pretty much is on top of things with the documents, whether or not they're being picked up by the market. The document does a really nice job, in each of the five areas of explaining the area itself. It's a nice high level description with the issues, threats, things to be worried about. The conclusion is on page 55. The remaining pages have tables of the standards and the analysis.

There was some push to address safety because much of the IoT implementation has to do with safety and privacy. Those concerns are mentioned on page two. Sometimes security standards can aid in limiting concerns, but the focus of the work group was not intended to

go in those areas. There is a reference to trustworthiness and it's not included. Cybersecurity is a big part of trustworthiness. A Venn diagram was used to determine the focus on privacy versus cybersecurity that came from the recent NIST Privacy document. For PII, there's an intersection between cybersecurity concerns and privacy concerns, but in no way is privacy a subset of cybersecurity or vice versa.

There are several legislative proposals, both Europe-wide and in individual nations. They are looking at creating duties of care and creating legal obligations for manufacturers and designers. From a Circular A-119 perspective, they could recognize those as voluntary consensus standards. The group tried to stay within the bounds of some of the federal policy in looking at what is a voluntary consensus standard, which was really defined by the organization that makes it. If it just looks functionally like a voluntary consensus body, then it was included. If it looked like it's a company effort, it was taken out.

This document's primary focus is to be useful to federal agencies as they look at IoT and cybersecurity standards and understand within their missions, how to make use of standards, where to engage what standards activities, as they go through their decision making. One of the reasons the document was put out into the public and private sector to get input is because the private sector knows about these standards and their adoption, and their opinion matters. That document is only a quality document if there is robust private sector input.

Extensive input and review is essential to almost everything, to have that kind of buy-in. The group hopes the Board members evangelize the importance of reading these documents and commenting on them, because that's what makes all of this work possible. The document will be a quality document, and the effort itself was worth doing. Just going through the effort of getting agencies to think about these things, provide input, meet and discuss them was very valuable.

Census 2020

Atri Kumar,U.S.Census Bureau

The Chair welcomed Mr. Atri Kumar, from the U.S. Census Bureau, to update the Board on the status of Census 2020. Mr. Kumar is the Chief of Decennial Information Technology Division at the Census Bureau. The Set Gap program, in the Decennial Directorate and managed by Mr. Kumar's division, is a program for data collection and processing. It is a group of solutions that the Bureau could use, instead of specifically creating something for 2020 Census.

The division went live with an internet self-response system. It is the final test before finalizing the solutions for 2020 and conduct the scalability testing on the solutions. End-to-end census testing is happening. It is collecting addresses through address canvassing. The responses then come through self-response, paper, internet, and questionnaire assistance and telephone centers.

They researched with technology including satellite imagery, it was established it would be helpful to local government by having a partnership with them. If entities provide reliable data on housing, it would cut back on the almost 600,000 address-listers that walked the streets 10 years ago. It was also found that the data could be collected by only having to visit 30 percent of

housing units. It means the other 70 percent of housing units can be identified through other technological means. It will reduce the number of address-listers to an estimated 75,000 to do that 30 percent of the housing unit address collection.

There is now an internet based self-response system, to introduce a "non-ID" means of collecting information. There is a chance that when the Census Bureau sends the census ID, that it could be lost in the mail. This type of data collection allows for individuals to provide an address and have the data map to the master address without the census ID.

Another innovative area is using administrative records and third-party data. Formerly, in order to identify vacant housing units, it was necessary to go and knock on doors. Now vacant housing units can be identified using existing administrative data the federal government already has.

The engineering field operations is another area where there is innovation. In 2010, information was collected on the housing units by mail carriers and enumerators knocked on the door, carried paper, and collected all this information. All that information had to be scanned into electronic processing systems. This time, there is a solution that will run on phones and the devices are being provisioned for the enumerators. The non-response follow up will be done using an application that runs on the phones.

Peak operations self-response involved live mailings requesting responses to the housing units in Providence, Rhode Island to actually go the website and respond. It served to validate the self-response contact strategy. They wanted to see if the numbers received really make sense.

When respondents have questions in relation to the questionnaire they receive or that they are using on the phone, the primary objective is to provide them help with the answers. If they agree to provide responses, the Bureau will accept them, but still use the internet self-response system to enter the data, meaning the operators are doing the work of the respondent.

The Bureau is embarking on a public campaign to let people know the census objective is to count only. The information will not be shared with other federal agencies. This count helps the communities and there is an effective way of making the case for, whether individuals are residents or non-residents, to respond. There will also be information at every Post Office. The effectiveness of the campaign really dictates as to how effective the census will be. Otherwise, people walking the streets and knocking on the doors as before.

For the 2020 census, the Bureau has implemented new solutions to prevent hacking attempts that have infiltrated its systems in the past. From a 2020 census perspective, the network has been segmented for the systems in the data center from the rest of the systems for the Census Bureau. Cloud systems have also been segmented. They will verify the information received is safe from penetration and have been working with DHS to complete penetration tests. They will be entering test readiness shortly, in May 2018.

Underwriters Laboratory IT Product Testing

Rachna Stegall, Underwriters Laboratory (UL)

Abel Torres, Underwriters Laboratory (UL)

Jeff Barksdale, Underwriters Laboratory (UL)

The Chair welcomed Ms. Rachna Stegall, Mr. Abel Torres, and Mr. Jeff Barksdale, all of UL to the meeting to brief the Board on Underwriters Laboratory IT Product Testing. The members of the Board introduced themselves to the speakers. Ms. Stegall is currently the Vice President of Connected Technologies at UL. Underwriters Laboratory has been around for over 120 years and its mission continues to be advancing safety. There are over 13,000 employees in about 100 countries helping manufacturers and asset owners get access to whatever countries that they want to sell their products to and ensure they bring safe and secure products to the marketplace. They also have a very mature security practice in the payment and the financial domain. They have over 450 security engineers around the world that specifically support the payment and financial domains.

The definition of safety has evolved to also security. The question asked by industry and governments is how to work with industry in developing a standards-based approach to evaluate cybersecurity for products, and provide some level of confidence in terms vulnerabilities, software weaknesses, and appropriate security performance.

On the industrial side, they work with industrial control manufacturers, renewables, which includes everything from wind and solar to power distribution to building automation types of products. They include energy management devices, chillers, furnaces, water heaters, lighting products, appliances and HVAC, consumer devices like cellphones and smart TVs, routers, servers, and automotive components. There is also a large practice to support customers in the healthcare field with medical devices and all the related regulatory requirements that come with that. There is a broad set of manufacturers that UL has worked with across those industries. In speaking to both the manufacturers, consumers, and governments, the message is clear that there are three key stakeholders looking for a level of support of the sector, not just from UL, but from the overall industry.

On the product manufacturer side, there are various reasons for looking for support. Some of them absolutely want a market differentiation. Some are because they have very limited technical expertise on the cybersecurity side. With a $500 million business, companies are telling us they have one IT security individual that is not only helping them secure their enterprise network, but is also helping advise their product and software development teams. They are looking for help with supplementing the knowledge gap and the resource gap in cybersecurity. There are multiple reasons for product manufacturers to need UL.

On the asset owners, retailers, and even governments' side, what's been asked is whether they are concerned about the risk. They would like to have some level of transparency that shows a common set of technical criteria that could be used to demonstrate that the supply chain has met that with some objective elements, so some level of transparency that has been validated.

UL has been working with insurance companies for years. UL started in1894 working with underwriters. Cyber insurance coverage is increasing faster than the understanding of aggregate losses. They are looking for capabilities and the tools to use to better lower risk and provide coverages.

Based on where each manufacturer related to their knowledge level, they need a different set of support as they increase that knowledge level. Some are trying to understand the first thing they can do to start increasing the product design to include secure design principles. Others are in need of some support and supplement knowledge in being able to evaluate a product. Lastly, there are definitely manufacturers that have good security practices for quite some time and they are looking for a way to differentiate their product in the marketplace.

UL started looking at the various standards and guidance and there are many great best practices that have been created by NIST, DHS, and various public-private consortiums. It can be very overwhelming for a manufacturer who does not understand what the first step is and at the same time, these documents have great best practices from the industry that can immediately be implemented.

UL helps its customers along that entire product development lifecycle. They like to get involved at the earliest stage and help manufacturers understand whether they are building products with secure and safety principles. Then, as it goes through that product development lifecycle, they can assess and verify that they have a vulnerability management process.

The European team is working through some options for privacy at this point. Devices are evaluated as to whether the data within the device is secure and how it actually communicates to other devices with what network protocols to set up. UL has not looked at an overall data privacy or privacy services right now.

UL has been thinking through the complexity of a physical mark for security or even a digital mark for cybersecurity. The biggest problem is the need to be able to help manufacturers in having assessable processes and audit to determine if they are maintaining that product until its end of life. Manufacturers are likely going to have a difficult time figuring out the processes, until they understand the secure design principles and actually build them into the product and test for it.

NIST Update

Matthew Scholl, NIST

Kevin Stine, NIST

Donna Dodson, NIST

The Chair welcomed Mr. Kevin Stine, Mr. Matthew Scholl, and Ms. Donna Dodson to the meeting to provide an update to the Board on NIST activities. Mr. Stine provided an update on Applied Cybersecurity Division activities across most of the key program areas. The division continues to seek opportunities to leverage applied aspects across all program areas, while trying to strengthen the relationships between many of the division's programs.

Mr. Barrett will speak in a lot more detail on the NIST Cybersecurity Framework. The second draft of version 1.1 was issued in December and there was a public comment period. Spring 2018 is still targeted as the final release of CSF version 1.1.

On the identity front, NIST issued a significant update to Publication 800-63-3, Digital

Identity Guideline. Identity is a topic that spans both divisions within ITL, and other divisions as well. In IoT cybersecurity, Mr. Hogan and Ms. Carnahan talked about cybersecurity standards earlier today. Kat Megas leads that program.

On the NICE front, it's been very busy since the last Board meeting. In November 2017, NICE issued a draft NISTIR 8193 on the NICE framework work role capability indicators and have received a lot of public comments. The purpose of the NICE workforce framework document is to determine if a cybersecurity professional can perform the work roles defined within the cybersecurity workforce framework. It provides examples of capability indicators, potential recommended education, certification, training, experiential learning, and other types of criteria that could signal an increased ability for an individual to perform particular work roles and capabilities. It is out in draft, feedback continues and we anticipate finalizing it this fiscal year.

In November, the seventh annual NICE conference was held in Dayton, OH. It was well attended. The 2018 conference is in planning. The first annual National Cybersecurity Career Awareness Week was held November 13-18. It was very much viewed as a campaign to focus on local, regional, national, and possibly international interest in inspiring and engaging people that may be interested in cybersecurity careers. Currently, the 2018 event is scheduled for November 12th - 17th. More information will come out in the coming weeks and months.

NIST hosted the NICE K-12 Cybersecurity Education Conference in Nashville, TN. There were a lot of educators, and several panels of high school students taking cybersecurity or related discipline classes. The 2018 conference is being planned for early December in San Antonio, TX.

There are plans to announce the addition of an apprenticeship subgroup to the NICE working groups. There are five subgroups: K-12, Collegiate, Training and Certifications, Competitions, and Workforce Management. There is a significant amount of interest in apprenticeships and a lot of activity on the industry side as well as interest on the government side. The NICE working group structure has been productive and valuable in many of these other areas. The group will be formulated soon and more will be available shortly.

There's a lot of joint work happening with the FISMA team relating to further integration of privacy controls into the NIST SP 800-53 control catalog. Publication 800-37 Rev. 2 is expecting an initial public draft in May, 2018, with a final public draft due in July, with final publication in October. A Federal Register notice will be coming out asking for comments and inputs to FIPS 200. It is expected before the end of March or early April. The goal is to publicly post that request for feedback. The Publication 800- 53 Rev 5 is set for final public draft in October, 2018. There will be comment on that draft prior to the initial public draft of March, 2019.

An informative mapping of the 800-53 privacy controls is being planned including the General Data Protection Regulation (GDPR). There's some informative mapping on how the different requirements within GDPR and how the different controls within 800-53 could relate from a privacy perspective. It's still being determined whether that's going to be an

appendix within 800-53, in a control catalogue, or as a companion resource. There have been requests for this information.

There will be a privacy roundtable in Washington, D.C. on March 29th. It will be associated with the IAPP professional summit that week. On the NCCoE front, in January a final version of publication 1800-6 on DNS-based email security was released. It describes a security platform for trustworthy email exchanges across organizational boundaries, by taking advantage of the capabilities of DNS.

There are a couple of additional projects that are coming up soon. A final practice guide was issued with a rough project description on financial sector asset management. The Federal Register notice is working its way through the process to encourage industry participation.

Mr. Stine was at HIMSS18 in March for a few days. There was a lot of excitement around the work NIST did related to wireless infusion pump security. There is a lot of interest and excitement over the next project regarding securing the picture archive and communication systems. A Federal Register notice to seek industry interest and participation in supporting these projects is coming. The project entails securing the picture archive in medical communication systems. It is the system that takes images from MRIs and CAT Scans and X-Rays and provides archiving and storage and access to those images. It has a lot of sensitive information, many images and data that would be attractive to other people. It was found approximately 50 percent of deployed medical devices were infusion pumps. It aligned nicely with the first medical device project that NIST worked on. There is a new project on the certificate management side. The second project that's ready to start moving is the IoT project to implement some security protocols from both ISAO and IETF to really minimize the loop around botnet.

There are four big buckets of activity: cryptography, testing and automation, risk management work and some research work.

This April, NIST is holding the First PQC Standardization Conference in Florida, where all Round 1 candidates will come to present their algorithms to the masses. It's their opportunity to stand up, explain their rationale, design reasoning, threat models and get direct feedback from the community. Once that is completed, they will go into round two, where other aspects of the algorithms will be looked at for selection. Things related to performance, retention requirements, size and space used in chips and gigs. This first round is really about the easy breaks. The second round, the cryptanalysis of the individual submission will be examined, as well as the performance characteristics..

A blockchain document, "Foundations of Blockchain," has been published to explain some of the underlying elements that make up a blockchain, how they can be used in different settings, both permission and permission-less blockchains, external blockchains, internal blockchains and then some of the underlying cryptography that a blockchain uses that NIST would consider to be sound. It really is a primer to explain what a blockchain is, and some use cases for blockchains.

The next step in the blockchain research is to look at things like specific blockchain use case implementations and extensions of how blockchain can be used. In some of the

models, a distributed block chain requiring proof of work is not efficient. But there are times when there is a non-distributed blockchain and permission blockchain where it's constrained to an enterprise or an organization. The proof of work requirement is either much less or it's done instead through that central authority. There's been a number of articles recently about blockchain or electric technology and how it impacts privacy because certain aspects of blockchain actually can be very protective of privacy and other aspects of blockchain technology can be quite problematic.

NIST will start aggressively developing a lightweight cryptographic algorithm, meaning one that will work in IoT constrained spaces. It must be a crypto-algorithm that still has cryptographic strong properties, but can operate in small, micro-controller devices, very low memory storage and very little bandwidth and/or combinations of those three aspects.

There's a public standards search going on right now called the CAESAR competition, which is also looking at lightweight cryptography and the eighth CAESAR competition just finished its public round one, so NIST is going to utilize and work with them as much as possible on lightweight cryptography. In the cryptographic maintenance area, they are looking at what algorithms are now legacy that potentially can be deprecated to avoid things getting out of hand. On the research side, in testing and automation, the IBM folks who talked to us about AI are folks who are working at NIST in looking at how to use machine learning and artificial intelligence in the National Vulnerability Database.

NIST is working with industry and standards bodies on a standard that uses a software ID tag, S-W-I-D or SWID. The goal is to have industry build software ID tags with their libraries and with their software, so that this software delineation is super easy. It just comes with a SWID tag. Tags come back for inspection with a system request.

The other thing for long-term is to tag libraries that have vulnerabilities because sometimes people just reach out and reuse code and reuse libraries. Vulnerabilities re-emerge, come back into systems, and come back into infrastructures through library and bad code readings. It isn't known specifically what code is tagged with what vulnerability.

In risk management, some of the risk management framework (RMF) documents are coming out. First and foremost is integrating the cybersecurity framework into the federal risk management processes, where and how to overlay the cybersecurity framework as the enterprise, expression, and understanding of risks, but use the processes of the risk management framework to generate agency and specific profiles of that framework.

They also want to incentivize agencies to shift resources away from low categorized systems to high category systems and HVAs. The question of how to incentivize the agencies was based on if there are eight lows and two highs, with an authority to operate (ATO) on those eight lows, the rating is eight out of ten. The goal is to be more nuanced and more specific to those assets and those systems that are rated high and high value. They are working with DHS and with ATC on what are the best ways to modify guidance to help agencies focus their resources, focus their time and put a priority on those high value assets.

The RMF was designed to be applied to an individual system and system-by-system. What's happening now is there's a more enterprise view of prioritizing on those individual system

RMF applications. NIST wants to incentivize that Enterprise view in a much more aggressive manner. In the new version of 800-37, it constitutes a shift in the RMF process to think about the ramifications of a low rating level, and consider the feasibility of outsourcing to a federal cloud provider. There is a distinction between enterprise level risk assessment and network or system assessments. They are looking for mechanisms and methods so that agencies can have a reasonable, informed and sound way to scope down their baselines, and apply limited resources to things that matter. The goal is to bring this GOVCAR toolset, these concepts into future updates to risk assessments to say that these are mechanisms that agencies should use and then the output should put that into insuring that the focus is on the right things.

There is a long-term trend of interface abstraction. The future is going to be voice interfaced. There is interest in starting now along with the integration of voice, potential wearables and augmented reality, which might be the next mobile computer platform of the future. Today's tech interface rests in a single device on a phone. That interface might get spread out to one's persona to bring in more augmented reality, enable more voice and then have a richer interface with technology that brings its own security and privacy requirements, issues and thoughts as well.

DDoS Report for EO 13800 Update

Kevin Stine, NIST

Megan Doscher, NTIA

The Board Secretariat welcomed Kevin Stine of NIST and Megan Doscher of NTIA to update the Board on the DDoS Report for EO 13800. The DDoS report is expected to be delivered to the public on May11, 2018. There were two workshops and two requests for comments. Ms. Doscher will speak about what the comments and about what's the same and what's changing in the report. Mr. Stine will speak about the workshop.

The report is in the drafting phase right now and as a result, there are no final answers for a lot of things. Generally speaking, the six themes the draft started with remain the same. The five goals cited in the document remain the same. The first focuses on a sustainable marketplace. The second focuses on infrastructure. The third focuses more on the edge of the infrastructure. The fourth looks at building coalitions. The fifth concentrates on education and awareness. There are no big changes in those areas.

Some of the comments relate to things that are different. One thing that will be different is considering a new section with a conclusion and next steps. It came up often in the comments, looking for more concrete information on priorities and activities. There were a lot of questions such as, what is the U.S. government going to do to lead all these actions? The answer is the government is not planning to lead all the actions, because it would be inappropriate for the U.S. government to lead actions that are specific for different parts of industry to work together on a problem.

The government wants to be involved if it’s appropriate to be involved. There were four or five different comments that asked a question on government involvement. It was a global theme that carried over into the workshop as well questions on roles, for government, and industry, and what those roles are.

The question of who leads was talked about at the last workshop. There's some onus on industry to step up on some of the areas that have been identified in the past. What does the Board think the appropriate role for government would be? There is hope an independent entity agrees to coordinate. If nobody steps up, there may be a role for government in that case. The effort should not die because there's nobody driving the action. If the private sector doesn't step up and lead in the areas that are designated for it, then government does have a unique ability to convene industry and other stakeholders on certain issues.

Another frequently heard theme in the comments focused on how infrastructure was characterized throughout the report. The initial definition took a very broad and all-encompassing view of infrastructure. It wasn't really carried through the whole report. In some cases, the language is being broadened. In some cases, it is getting more specific. Sometimes "ISPs" were mentioned when infrastructure was intended. It's being modified in the report.

Comments were also received about international standards. It may not have been totally clear that the government always wants industry to take the lead. The report is being modified to be sure that it's very clear. There's no mention about foreign governments, or pushing forward a standard. Information sharing came up in a lot of different ways. It was assumed that "reporting" always meant instant reporting, but that was not intended. That is being clarified. There was some talk about automation pros and cons, some in the context of liability. There is a big box about liability in the report.

There will be more to describe the environment. Some of the commenters felt like a path or outcome was not fully expressed, but no specifics are being recommended yet. It's early at this stage to try to make a determination. The downside to any liability regulations is the potential to stifle innovation. If there were liability protections, there needs to be some sort of responsibilities that go along with those protections.

Discussion of individuals and small businesses came up a lot at the workshop. In the report, it was felt small business were considered too much as inadvertent perpetrators versus victims. That was a good thing to have pointed out. It is important to note small businesses are often victims, more so than large companies. That clarification is now included. Some comments talked about the idea that it could have been inferred from the report that the government didn't think individuals or small businesses really had any responsibility to secure their networks because it's not expected they will be technologically able to. It was not intended to state that. The language will be clarified further.

Several members of civil society were upset that they were not invited to participate in the report. That was inadvertent as well. The report is being modified to talk about successes civil society has had in the past and how they might be able to contribute going forward. In all the sessions of the workshop, there was always a small business concern or consideration.

It was certainly a recommendation that insufficient attention was paid to the needs of small businesses, both on the inadvertent perpetrator and the victim side. Small businesses make the technologies that are being used. Sometimes they're being used in ways that are

unintended.

There are also concerns if, for example, there is action on certification processes and related plans for devices, those can't be cost prohibitive to the point where small businesses won't be able to participate in that market activity. There are a lot of nuances and considerations where many of these potential recommendations or actions are concerned.

One of the things that was interesting in the workshop was the question of what can small businesses really do from a threats interpretation perspective when it takes financial firms three full-time people daily going through all of these things. There may not be an answer to that. The community doesn't have all the answers that are needed today. Certainly, there are actions that can be taken today, collectively or as individual organizations.

A workshop was hosted at the NCCoE from February 28th to March 1st, with roughly six hours of panel discussions and four hours of breakouts really to keep folks engaged. There was a lot of lively discussion and a lot of folks very engaged in the content. The workshop was a big success. One of the primary goals for this workshop was to get feedback and clarifications on the comments received through the comment period. The inputs received during the comment period were used to shape the agenda for the workshop.

The event started with a financial sector community panel. There have been many successes and great collaborations over the last several years specifically in these automated industry-rooted threats impacting financial infrastructure. Their perspective was extremely valuable. One of the interesting things in that panel was the focus on gaming. That was not included in the draft report. It played into a lot of the other discussions over the course of the day and was valuable.

Cybersecurity Framework version 1.1 Update

Matthew Barrett, NIST

The Board welcomed Mr. Matthew Barrett of NIST to the meeting to update the Board on the status of the Cybersecurity Framework version 1.1. The framework update has been going on for some time. The final framework will be out by the end of April, 2018. There were 86 comments from industry in the most recent round. Some commenters were trade associations representing large numbers of organizations across those 86 responses.

A lot of input on the draft roadmap was pretty critical infrastructure security oriented. It was published in December, 2017. More comments were actually received on the draft roadmap than on the framework itself. The transparent repeatable update process will also be released at the time the final of Framework 1.1 there was a gap between the original development of the framework and this one iteration, to make it a living document.

There were learned lessons along the way that will be channeled into a process to be published so all parties know what the repeatable process for framework updates look like in the future, including things like how often updates will occur and what will be in them. Those details will be forthcoming in April. They continue to focus programmatically on small-and medium business, on international alignments, on regulatory circumstances effective May, 2017 with Executive Order 13800.

The process has been occurring since the draft of NIST IR 8170. It was the framework implementation guidance for federal agencies intended to assist agency heads in responding to the implementation plan and the report owed to the President. There were eight proposed uses of the Cybersecurity Framework for federal organizations in that document. Those were uses seen in the private sector. NIST IR 8170 will be finalized in late spring 2018.

Internationally, there are a lot of exciting things happening. NIST is liaising with our colleagues at ANSI and they are in turn liaising to ISO IEC on relevant study periods and work products related to the cyber security framework. It includes things as foundational as debating, defining, and capturing the differences between information security and cyber security. It also includes an upcoming technical report 27103 which shows relationships between key ISO and IEC publications and the Cybersecurity Framework.

A comprehensive national framework might include coverage of identify, protect, detect, respond, recover. There are properties like these five areas from the cyber security framework that will be written into that technical specification. Mr. Barrett spent time reviewing Uruguay's adaptation of the cyber security framework. It was really interesting and clever. In some ways, it was similar to the Italian adaptation in that they have expressed priorities across the sub-categories and also have a maturity scale and qualifying criteria on a subcategory basis.

There are some important work products coming out of South America. It already reached version 3.1which was surprising. There is framework usage in Portugal. The Israelis have created an adaptation and a translation in Hebrew of cyber security framework and the Japanese translation always has been there.

NIST supported colleagues in Bermuda who are quite public about their use of the Cyber Security Framework, they use it alongside of SP 800-53 and the Risk Management Framework in their own governmental management of cyber security. They advocate these work products to their industries and NIST supported them with a workshop this past fall. They are proceeding with the concept of starter profiles that are subcategories most important to certain business functions, maybe business functions that are common amongst small business. That's something that's out on the horizon for late calendar year 2018 into early calendar year 2019 for that project.

The profile is the customization mechanism within the framework to define what subcategories are applicable, and what is more meaningful to a given circumstance. That's "profile" is preferred as a word rather than "light". They should be drawing on a common catalogue in their customizations in their profiles.

Collaborations have started with the Small Business Administration and trying to understand what are the most commonly occurring small businesses. They are talking about what their business functions look like and some of that's driven by government. Some of that is coming back through director participation and participation of other government colleagues in small business venues. It's at least starting there because there really isn't a good picture.

The work with the U.S. Coast Guard developing profiles for maritime has been completed.

There are three such profiles that exist for various maritime segments. In April, there is a co-hosted event at the Department of Commerce on behalf of the Financial Services Sector Coordinating Council and their financial services customization of Cybersecurity Framework. There will be a workshop to advance that work,

Upcoming plans include: April 2018, the final of Cybersecurity Framework version 1.1; in late spring 2018 finalizing NIST IR 8170; December 2018, a Spanish language translation of the framework. Six hundred sixty-eight billion dollars of the gross national product is produced by Spanish-speaking organizations. When version 1.1 is finalized, translation for the Spanish language version will start. The annual conference will be rebranding this year to the NIST Cybersecurity Risk Management Conference. It will be a three-day, three track conference coming up in late summer to fall. Then in the late fall into winter the starter profile concept will be more the center of attention.

Board Review and Discussion

The Board discussed the following areas of interest concerning letters, topics of interest, and future meetings.

Future Meetings:

1. The next ISPAB meeting tentatively scheduled for June, during the week of the 17th.
    a. Board members are suggested to report their summer and vacation schedules so that a consensus can be met as to when the next meeting will occur.
    b. The meeting cannot happen the week prior as the ITL laboratory has a meeting scheduled with the National Research Council.
    c. The location of the meeting is TBD, but there seems to be a preference between the Access Board and American University in Washington.
2. The meeting minutes will be sent out to the Board for review within the next couple of weeks.
3. The Board agreed the two-day schedule of longer days is more acceptable than the three day schedule that was done in the past.
    a. This seems to work fine as it allows at least an hour for presentations and discussions and some tend to end earlier.
    b. Due to the short discussions during some presentations, there is a tendency to end a lot earlier when to the meeting lasts three days.

Future Meeting Topics:

1. There has been discussion about a follow-up for de-identification for future topic.
2. A discussion on supply chain and workforce.
3. A deeper dive into a discussion on Blockchain.
4. Artificial Intelligence (AI) and the security surrounding it.
5. Looking into social engineering awareness and education programs.
6. Discussions with the IRS and FTC
7. We are due for another legislative round up.

8. Have a specific presentation on the new draft of SP 800-37, as it would have just come out.
9. Invite Rob Joyce again to discuss critical infrastructure adoption, framework, and
10. The small business profile and getting it out to the small businesses. It would be nice to have a little more granularity on how well they are doing measuring.
11. Given the executive order and all the due-outs from that, it is time to get a briefing from the White House.
12. It would not be surprising if there was another executive order that came out between now and June, so, we should hold a spot or two in case that happens.

Areas of Interest:

1. There is already been a lot of attention put on PCLOB, whether or not the Board have any role to play may be worth discussing.

Meeting Recessed

The meeting recessed at 3:17 p.m., Eastern Time.

List of Attendees

Last Name	First Name	Affiliation	Role
Scholl	Matt	NIST	DFO / Presenter
Barksdale	Jeff	UL	Presenter
Barrett	Matthew	NIST	Presenter
Carnahan	Lisa	NIST	Presenter
Connelly	Sean	DHS	Presenter
Dodson	Donna	NIST	Presenter
Doscher	Megan	NTIA	Presenter
Gilmore	Mike	GAO	Presenter
Hogan	Michael	NIST	Presenter
Jackson	Crystal	DHS	Presenter
Kumar	Atri	US Census Bureau	Presenter
Malhotra	Kush	GAO	Presenter
Marrios	Nick	GAO	Presenter
Ramamurthy	Karthi	IBM	Presenter
Romine	Charles, Dr.	NIST	Presenter
Stegall	Rachna	UL	Presenter
Stieren	David	NIST	Presenter
Stine	Kevin	NIST	Presenter
Suagee	Jason	NIST	Presenter
Torres	Abel	UL	Presenter
Drake	Robin	Exeter Government Services	Staff
Salisbury	Warren	Exeter Government Services	Staff
Gaurav	Pal	StackArmor	Visitor
Heyman	Mat	Impresa Solutions	Visitor

Last Name	First Name	Affiliation	Role
Kerban	Jason	Department of State	Visitor
Mohindru	Aman	IBM	Visitor
Myles-Brown	Kathy	G2	Visitor
Nong	Nai	DHS	Visitor
Savickis	Mari	CHIME/AEHIS	Visitor
Seaborn	Lydia	DHS	Visitor
Senholzi	Peter	G2	Visitor
Sokol NIST	Annie	NIST	Visitor
St. Pierre	Jim	NIST	Visitor
Miller	Jason	FedNews	Visitor/Media