INFORMATION SECURITY AND PRIVACY ADVISORY BOARD

Established by the Computer Security Act of 1987
[Amended by the Federal Information Security Management Act of 2002]

MEETING MINUTES

June 15, 16, and 17, 2016

U. S. Access Board
1331 F Street N.W. Suite 800, Washington, DC, 20004, (202) 898-4000

Board Members	Board Secretariat and NIST staff
Chris Boyer, AT&T, Chair, ISPAB Dave Cullinane, TruSTAR Technologies (joined via phone) Greg Garcia, McBee Strategics Consulting Jeffery Greene, Esq., Symantec Corporation Toby Levin, Retired Edward Roback, US Department of Treasury Gale S. Stone, Social Security Administration J. Daniel Toler, US Department of Homeland Security **Absent with Regrets:** Dr. Ana I. Anton, Georgia Institute of Technology John R. Centafont, NSA	Matt Scholl, NIST Annie Sokol, DFO, NIST Donna Kimball, NIST Robin Drake, Exeter Government Services, LLC Warren Salisbury, Exeter Government Services, LLC

*** Footnotes are added to provide relevant or additional information*

Wednesday, June 15, 2016

The Chair opened the meeting at 8:36 a.m.

Welcome and Remarks

Chris Boyer, Chair, ISPAB, Assistant Vice President, Global Public Policy, AT&T

Mr. Boyer opened the meeting with remarks on the future of the Information Security and Privacy Advisory Board (ISPAB), and expressed the desire to have Dr. Charles Romine open future Board meetings by discussing new priorities for the National Institute of Standards and Technology (NIST) and what the agency is doing, so that the Board can prioritize issues and provide feedback to NIST on those that are the most relevant. Mr. Boyer expressed the desire to work more on letters or white papers on particular topics that represent major issues. The internet of things may be one topic area. The Board should give some thought to additional topic areas.

The Chair asked the Board members to introduce themselves and provide an update if there was one: Ms. Toby Levin has been active with security issues. Mr. Jeff Greene has been looking at Internet of Things (IoT) issues in

automotive and healthcare, and security and privacy issues. Mr. Ed Roback with the US Department of Treasury, has been dealing with infrastructure issues in the financial sector. He is involved with two task forces related to payment security.

The Chair noted the National Cybersecurity Incident Response Plan would be out in October, 2016. There is a lot of development work going on, and just this week DHS kicked off an effort to develop it and asked for representatives from critical infrastructure. Mr. Boyer was selected as one of the communications representatives. On June 2nd, the Homeland Security Advisory Council (HSAC) made some responses regarding incident response. The financial sector, communication sector, and energy sector all provided recommendations on how the private sector should organize itself in the event of a major cybersecurity event. There is now a tri-sector working group representing communications, energy, and finance as a result to try to develop a playbook on how things should work in the event of a major cascading cyber-attack. We hope that it never happens, but it is best to be prepared. HSAC is concerned with how things should work in the event of an attack. It is an important initiative, and worth looking at in October when there is something to review.

Presentation on Secure Engineering and Cybersecurity Resilience [1]

Dr. Ron Ross, Fellow, NIST

Presentation – Progress Update: Federal Cybersecurity [2]

The engineering project is now four years old. We are getting close to finalizing in 2016. This engineering work is going to be fundamentally important to the cybersecurity work we are trying to achieve. The Office of Personnel Management (OPM) taught us a lot of lessons and had a personal impact for many. It is an inter-generational breach because family information was also involved. The information is out, and it can never be gotten back. The first breach involved federal employees, especially those with clearances. The second one Mr. Ross was involved in was the Anthem Blue Cross Blue Shield (BCBS) breach. The Ashley Madison breach, along with these others, exposed a massive amount of information. The attackers were looking for a cross section of federal employees with top secret clearances who might also have potentially embarrassing issues they would want to keep quiet. This breach goes well beyond credit protection, but actually enters the world of espionage as it involves other personal issues. These are the kind of things that make us want to be as secure as we can be. Technology is moving at warp speed, and we are doing our best to hang on. Technology is exceeding the capability to protect it. There is a cultural issue. The U.S. has the best technology in the world, and we buy a lot of it. All of the new technology means additional complexity. It also means that attacks will increase, accordingly. It is pretty much uncontrolled.

There are about six billion devices making up the internet of things, and five million more are being added every day. We need to try to figure out how to advance the technology, and keep risk at a manageable level. Risk levels will vary by sectors. Each must determine what level of risk is acceptable.

1 DRAFT NIST SP 800-184 Guide for Cybersecurity Event Recovery http://csrc.nist.gov/publications/drafts/800-184/sp800_184_draft.pdf
2 http://csrc.nist.gov/groups/SMA/ispab/documents/minutes/2016-06/1_secure-engineering_rross.pdf

The attacks on the World Trade Center were kinetic attacks. A kinetic attack can be felt. Cyber-attacks are different from kinetic attacks. They exist below the radar. Effects of cyber-attacks may not be felt for some time. The average time for a cyber-attack to be detected is 200 days. The OPM breach was not detected for seven months. A lot of damage can be done in that time. We are trying to look at the problem form a different perspective. Even after four decades in cybersecurity, there are still breaches and damage going on. Cyber-hygiene is doing a lot of good, but we're not doing as well "below the water line". The problem we're facing today is there are a growing number of unknown vulnerabilities. Vulnerabilities are growing, and the unknown is expanding exponentially.

The only way to control the growth is through good architecture and engineering techniques. We will never be able to stop all attacks. We are drowning in threat intelligence and information because we are collecting lots of data. The only way to look at it long term is to consider what happens if we can't stop the adversary, we can limit the damage. If the damage at OPM could have been controlled, possibly a much smaller percentage of personal information may have gotten out.

The first thing is to harden the target. It makes the system more penetration resistant, but it's not perfect. Firewalls, encryption, two-factor authentication are a means to this end. Measures can make systems resistant to attacks, but there will still be successes. Limiting the damage is key. There are two approaches. Limiting damage by system engineering is one way. We can limit the time the adversary has in the system.

The other way is to limit horizontal movement across systems. System resilience needs to increase. Work will be ongoing to apply the principles of security engineering to improve systems. The intelligence community has introduced the High Value Asset Initiative, which looks at system vulnerability from the adversary's perspective. The ultimate goal is to make the system more trustworthy, resilient, maintainable, and survivable through innovative security engineering. Then, by doing diligent risk assessment, an organization can better determine the cost it's willing to invest to make the system as trustworthy as it needs to be for that particular organization.

There are overlaps between safety and security. These overlaps can be utilized. Risk assessments are still important by studying the approaches taken by the National Aeronautics and Space Administration (NASA) and the Federal Aviation Administration (FAA). The project found a great deal of overlap in areas of vulnerability and security and decided against making NIST Special Publication (SP) 800-160 a security requirement, but rather to investigate the best systems, software, and processes and to adopt those strategies.

Risk assessment is presented as comprising four components: threat, vulnerability, impact, and likelihood. Cyber-security may thus be understood from the perspective of any of the four components. In SP 800-160, the project proposes starting from the perspective of impact because it provides the best basis for understanding the value of an organization's assets. It also provides the base on which to "engineer up": to identify high-value assets, and then build and strengthen the domain accordingly.

Dr. Ross noted the SP 800-160 [3] is probably the most important document he has ever worked on because it addresses cyber-security problems "below the water line" and presents the need to develop command and control capabilities to better understand the complex environment into which security strategies must be implemented.

There are two groups of processes (non-technical and technical) as mandated by ISO/IEC/IEEE 1528:2015 as the most robust standard studied by the project, since it brought the security team into the world of engineering solutions. Permission was granted by the International Organization for Standardization (ISO), International Electrotechnical Commission (IEC), and the Institute of Electrical and Electronics Engineers (IEEE) to Dr. Ross's team to implement their joint set of standards.

The non-technical processes of the ISO/IEC/IEEE 15288:2015 *Systems and software engineering – System life cycle processes* include Project Planning, Project Assessment and Control; management specific to Decision-Making, Risk, Configuration, Information, Life Cycle Model, Infrastructure, Portfolio, Human Resources, Quality, and Knowledge; Measurement, Quality Assurance, and Acquisition and Supply.

The technical processes include Business or Mission Analysis, Stakeholder Needs and Requirements Definition, System Requirements Definition, Architecture Definition, Design Definition, System Analysis, Implementation, Integration, Verification, Transition, Validation, Operation, Maintenance, and Disposal.

The IEEE and ISO updated standards last year. "Business or mission analysis" and "Stakeholder needs and requirements definition" are new areas. The object is to have the mission business owners take responsibility for the security of the project. Stakeholder needs will drive stakeholder protection needs. Those requirements get built into the process as system security requirements, which then get built into the process. It carries through the entire Standards Development Life Cycle (SDLC).

Dr. Ross presented an organization's purpose statements as examples of built-out cyber-security activities specific to the Human Resource Management Process from the perspectives of systems engineering and systems security engineering. The first task took the form of a gap analysis to identify which people are needed to implement the standard and which people are currently available. The gap analysis would drive the decision (the "outcome") either to hire additional staff or to train existing employees.

As with all other processes, systems engineering and systems security engineering human resource processes are ultimately linked to the corresponding process as defined in ISO/IEC/IEEE 15288:2015, and also linked to related publications addressing the specific process. This effort has the advantage of demonstrating that an organization's processes might already be in place and provide considerations to achieve the standard. These processes sit on top of all the other work that's been done.

For every process, there is an outcome. The group developed security based outcomes based on ISO. There are high level activities and security tasks to complete those activities. The group is avoiding the word "guidance" in favor of the word "considerations" to avoid agencies from using of NIST documents as compliance weapons in issuing RFPs, etc.

---

3 NIST Special Publication 800-160 (2nd Public Draft) Systems Security Engineering – Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems
http://csrc.nist.gov/publications/drafts/800-160/sp800_160_second-draft.pdf

Is intended to be used as a process that draws on the work that has been done to date. It does not replace SP 800-53. There are three target audiences: government, industry, and academia. This work will be the basis of the National Security Agency's (NSA) new security engineering program that will be pushed out to 80 colleges and universities.

The program needs to be facilitated by the Chief Information Security Officers (CISOs). Everything great that happens in this country is based on government and private sector working together. We need to get people who are already engaged, to incorporate this process and continue working. SP 800-160 DRAFT will be the flagship document with other appendices being broken into separate publications.

An incredible amount of money was spent on the cold war and mutually assured destruction. It was probably one of the most expensive ever made. But the assets were the highest level of protection. Clean up of attacks will always be an order of magnitude more than the cost of protection.

Ms. Levin: This is a tool for leadership, governance, and accountability. It allows for cybersecurity to be baked in to all processes. There needs to be structured detail to meet the challenge. It supports all development processes. Should we focus people more toward a fundamental change of architecture, rather than information sharing?

Dr. Ross: Some things tend to absorb too many clock cycles that could be used to focus on major problems. Future Board discussion on prevention and focus on threats.

Updates from Office of Management and Budget (OMB)

Trevor Rudolph, Chief, Cyber and National Security Unit, Office of the Federal CIO, Office of Management and Budget (OMB), The White House

Mr. Rudolph is providing an update on the progress in cybersecurity. Some significant events have happened since this time last year. Are we actually making progress? Headlines may indicate the answer is no. The public narrative has been pretty bleak as a whole but is this the whole story? It is difficult to measure how many incidents have been prevented by good security. Government is not the only one battling threats. Governments have to deal with issues others do not.

There are indicators that progress is being made, even if the pace is not what everyone wants. Most threats are due to poor authentication and unpatched vulnerabilities. Major progress in implementing the PIV card happened during the cybersecurity sprint. We need to be close to 100% implementation. It is an indication that the numbers are trending in the right direction. There has been a marked decrease in authentication incidents following the sprint.

There has been a 99% reduction in critical vulnerabilities since the sprint. There are still challenges to progress in legacy IT, fragmented management of IT, and workforce challenges. Every agency makes its own risk decisions, hundreds every year. It is not a sustainable position for risk management.

Mr. Greene: What do you think caused the uptick, and how to we keep the level of awareness that exists now from dropping off in the future?

Mr. Rudolph: There is leadership engagement that exists now that has not been evident in the past. Staff execution has been critical in the progress that has been made. The OMB had the resources for the first time to conduct stringent oversight. Are adversaries still banging on the same doors? We are making it more difficult. Advanced threat actors will not be deterred for long.

The government is becoming more agile. The next presidential transition will be critical. A lot of resources are being dedicated to the transition teams. It must be a priority on day one and having people power ready to work. The intent of the cyber sprint was to implement rapid improvement of cyber hygiene. Cybersecurity Information Sharing Act (CISA) intended to tighten policies.

Much more energy and leadership are now being focused on Cybersecurity National Action Plan (CNAP) in order to have something more codified by the end of the current administration. It lists $19B in proposed cybersecurity spending. The White House must be regularly engaged in this process.

Key federal initiatives include securing high value assets, the IT modernization fund, centralized IT services for small agencies, cybersecurity workforce strategy, and cyber defense teams. Who provides the service and who owns the risk? If we can work out what's right for small agencies, it will have big implications for the entire Federal Government. OPM will soon release the first cybersecurity workforce strategy for the entire Federal Government. Cyber defense teams working out of Department of Homeland Security (DHS) are being considered. We will also be working with industry to supplement this work.

The challenge moving forward is prioritization, codification, implementation (quick results), and communications. The government and private sector are different. How are things being adjusted? We are consulting with government experts and the private sector. A best practice in one organization may not be a best practice for all organizations.

The Chair: The Chair expressed appreciation for the information being presented. Can agencies avail themselves of these services in protecting high value assets?

Mr. Rudolph: It is possible.

The Chair: Can external providers assist small agencies?

Mr. Rudolph: Yes, that is possible.

Executive Order – Establishment of the Federal Privacy Council [4]

Marc Groman, Senior Advisor for Privacy, The White House

Mr. Groman fills a new role in the Executive Office of the President. He works with agencies like OMB and others, bridging the gap between law, policy, and privacy. One of the first tasks was to evaluate the current state of play. The definition of privacy covers intelligence community issues, consumer issues, and all other privacy issues related to the Federal Government. It covers people's privacy relating to collection, creation, management, storage, and transmitting information about people.

Privacy and security are not the same thing. It is not sufficient for the government to focus only on security programs and not focus on privacy. Security is about the integrity and protection of data. Privacy involves a far broader range of issues. Can goals be achieved with less information or less sensitive information? How information is used is critical. Uses of information must not contradict policies or intentions for use across agencies. Privacy questions must be considered at the beginning of any project. Privacy is about data governance. If we don't know what's happening with the data we have, it's very hard to protect.

The support from the current administration has been extraordinary. Mr. Groman's work centers on two issues. We are looking to transform how information is used. We want to move away from the perception that privacy is an exercise. Privacy must be comprehensive, department-wide, strategic, and part of an enterprise-wide risk based program. We must get away from the check-the-box notion and move to something more encompassing. It involves focus on people and focus on policy. Both are critical and one is no good without the other.

An executive order was given during the CNAP rollout for the first ever federal privacy council. Privacy is more important in a digital world where more is being shared than ever before. The council is fully active now. One of the top priorities is the federal workforce. There is a perception that privacy impedes innovation. It is false. Privacy laws are not a roadblock, but a promoter of innovation. Privacy will allow for large scale adoption of innovation. Privacy helps ensure long term success, and create trust with the American people.

A competency model for privacy is being developed. Job descriptions are being developed. Possibly a job series will be developed in the long term. There are two tracks in privacy education. There is a technology track for privacy professionals. Privacy professionals need to understand technology in order to do their jobs. The council now has funding from GSA. Leadership is critical to any program. The leader sets the tone for any project. Agencies are now posting for and filling Chief Privacy Officer positions. There is a commitment across the government to these privacy initiatives.

Privacy programs are situated in many different places across agencies. It can be in the CIO or General Counsel's office. The privacy act appendix of the Circular No. A-130 Revised [5] will become its own circular, the A-108.

---

4 https://www.whitehouse.gov/the-press-office/2016/02/09/executive-order-establishment-federal-privacy-council; https://www.whitehouse.gov/blog/2015/12/01/prepared-remarks-omb-director-shaun-donovan-federal-privacy-summit

5 https://www.whitehouse.gov/omb/circulars_a130

Mr. Groman's office is doing an overhaul of a document that details responses to incidents involving personally identifiable information. They are working with the Office of the Federal CIO (CIO) and others. Incident responses always involve forensics. Once it's known the incidents involve Personally Identifiable Information (PII), it's not just a cybersecurity response. The document will have an opportunity for public comment. They are working with the Federal Trade Commission and others.

Ms. Levin: Does the council review System of Records Notices (SORNs)?

Mr. Groman: They are reviewed, and I may be part of that process, if appropriate. Privacy act guidance (circular A-108) is in process, and should be processed before the current administration ends.

Training is available that can be used by other agencies. We are looking at how to provide the best training for free. There is a budget for the council itself, but not for training. There was a privacy council summit last year.

Having privacy involved at the beginning, is very valuable. It needs to be a focus point going forward. NIST is thinking about privacy-preserving technologies. NIST researchers have identified close to 50 public privacy incidents that are not connected to any cybersecurity incident.

The Chair: We will want to consider how we can be of assistance to Mr. Groman's effort. Will the risk management framework be incorporated into privacy impact assessments? Yes, but the time depends on how one defines the time-horizon.

Commission on Enhancing National Cybersecurity [6] [7]

Kiersten Todt, Executive Director, Commission on Enhancing National Cybersecurity

Kevin Stine, Chief, Applied Cybersecurity Division ITL, NIST

There is new information to report since the last time the Board heard about the commission. The commissioners were officially announced in April following the previous ISPAB meeting. There are eight (8) areas of engagement identified in the Executive Order - Commission on Enhancing National Cybersecurity. Three additional areas also identified by commissioners as being relevant. All of these topics fall within the digital economy.

The audience for the commission ranges from individual consumer to state, local, tribal, and Federal Governments. The final report is due to the President on December 1, 2016. Five workshops are being held around the country: New York, NY, was held in May; June was held in Berkeley, CA; followed by Houston, TX, in July; and Minneapolis, MN, in August. The final workshop will be in Washington, D.C. in September dealing with governance.

6 NIST Cybersecurity Commission http://www.nist.gov/cybercommission/commission-meetings.cfm
7 Meeting Minutes, April 14, 2016 http://www.nist.gov/cybercommission/upload/Meeting_Minutes_April_14.pdf

A request for information (RFI) will be out to the public in two weeks. Working groups have been created for each of the subject areas mentioned above. They will be writing draft recommendations that ultimately will feed into the report.

The types of cybersecurity problems are not changing, but we continue to have the same issues. We need to understand the people element: what the incentives are for cybersecurity, and what motivates them. Is cyber security perceived as a burden on the individual? There needs to be a national cybersecurity education campaign. It needs to be integrated into the culture. We have been asked to address critical infrastructure. How is this defined in terms of the digital economy? Secretary Penny Pritzker and the Department of Commerce (DOC) have started to work on defining this area.

In the report and recommendations there will be three (3) categories: New Ideas, a Collection of Best Practices and Lessons Learned, and Ideas/Strategies that have not been implemented yet. The work of the National Initiative for Cybersecurity Education (NICE) is integral to the effort.

We have two calls of action to the ISPAB: The commission is interested in considerations of what we should be looking at. The RFI is also a vehicle for gathering these considerations. We are also looking at concrete and actionable actions that can be undertaken along with long-term plans for actions in the future. The commission is looking for input from the Board on how to socialize contents of the report, and input on what's been done and not worked, previously.

The burden for safety cannot be placed on the consumer. What we came to understand is that it must be baked in. Safety features are built into cars; it's not a choice. Devices must come with the protections built in also. There is a global entrepreneurship summit in CA next week. Entrepreneurs have the opportunity to create security. Cybersecurity is not the barrier to innovation that it is perceived to be. The goal is to make doing the right thing easy, and the wrong thing, hard.

The incentive construct might be better served to be perceived as a basic security rather than as has been previously portrayed. Is the end user really able to protect themselves even if they are aware? It may be better to remove the end user from the equation of protection. For companies, it may be expedient to limit what users can and cannot do. We can't expect the government to be solely responsible for an education/awareness campaign. There must be collaboration among all parties involved. How can we make raising awareness efforts more effective? We must also maintain awareness of technology changes going on presently. It is a challenge to communicate with all employees in a companies to educate them.

Higher awareness can create demand for security. Targeting awareness of the total cost of ownership in products that are less secure. It means higher costs for breaches, etc. Establishing a common bar for insurers, so that insurance isn't issued if security does not meet the set bar. There are a number of tie-in areas. Devices can be certified. The supply chain is directly related.

We have to be careful what certifications mean. People may trust certifications too much. There are still unknown certifications and zero days to come. It's not just devices, it's also browsers. Browsers need to have the proper defaults and messaging to users. Internet Service Providers (ISPs) and browsers should be working together to

detect presence of firewalls, etc. What do hacks for certified devices mean? It depends on what's been certified. It comes back to what is defined to be user error, etc.

If the number of devices to be secured is reduced, security issues may be reduced, thereby reducing the burden on the consumer as much as possible. To the extent possible, security is not a choice.

Mobility is something the commission is aware of. The scope it will play in the final report has not been determined. The commission is examining the digital economy in five or ten years. We may need to stop calling our mobile phones, phones.

Cybercommission@nist.gov is one venue for communications.

Board action: Letter to the Chair and Co-Chair also possible regarding recommendations. Mr. Boyer will try to determine a process for a letter. Letter in the next 30 days. Possibly an open call also.

Board: We should have Peiter C. Zatko, better known as Mudge, come to the next ISPAB meeting.

Cybersecurity Information Sharing Act (CISA) Implementation

Jocelyn Aqua, Senior Component Official for Privacy/Counsel for Law and Policy, National Security Division, U.S. Department of Justice (DOJ)

Leonard Bailey, Senior Counsel, U.S. DOJ

Matt Shabat, J.D., Strategist and Performance Manager, Office of Cybersecurity and Communications, U.S. Department of Homeland Security

CISA is intended to give new authorities to private agencies. It addresses complexities of certain standards that have hindered people in the past. It monitors information in information systems, apply defensive measures, and define threats.

CISA provides the authority to monitor information or information systems. It overcomes any conflicting authority. CISA authorizes communications. Authority is bound to cybersecurity threats or vulnerabilities. There is monitoring protection provided in CISA.

Application of defensive measures are defined under the statutes. Defensive measures are defined to "exclude access to, destroy, or substantially harm" another system. Defensive measures can be allowed to another system owner's system. Owners are allowed to protect their rights and property.

Mr. Roback: The statute offers no guidance on the question.

Ms. Aqua: It will discuss the privacy and civil liberty guidelines built into CISA. We met with government and privacy advocates outside the government. They cover retention, receipt, etc., of threat indicators. The standard for reviewing cyber threat indicators got a lot of attention. It is focused on information deemed to be directly

relevant to the government or others. It can be for law enforcement purpose or national security. The main difference in the guidelines is there is more clarity. The portal takes in information from other sources.

Matt Shabat, J.D., Strategist and Performance Manager, Office of Cybersecurity and Communications, U.S. Department of Homeland Security

It is important to understand that sharing is consistent with the act. Information is scrubbed, there is no personally identifiable information (PII). There are protections in place during sharing by federal and non-federal agencies. State and local governments are treated somewhat differently. Sharing in electronic format is what is focused on. DHS has an obligation to share threat indicators received from non-federal entities with all parties.

An automated scrub is applied, and human analysis used if necessary before CISA passed the revised profile. Some privacy protections were built in to the profile. In the future, we would like to be able to add reputation scores, sharer/sharing discussion and related protections. The basis of the statute is to only share directly relevant information. Existing sharing relationships can continue to operate and fall with within protections. A lot of sharing outside of liability protection needs to be recognized. The concern is that there may end up being no sharing. People are waiting to see what the final guidelines will look like.

Updates on NIST Post-Quantum Cryptography Standardization Plan [8] [9]

Lily Chen, Ph.D., Acting Group Manager, Computer Security Division (CSD), Information Technology Laboratory (ITL), NIST

Dustin Moody, Ph.D., Computer Scientist, CSD, ITL, NIST

Quantum computers are different from classical computers. They operate on quantum mechanics. A qubit can hold 2 to the n bits at the same time. If a large scale quantum computer would be built, it would have a large impact on cryptography. Standards that would be at risk at NIST include – FIPS 186 and SP-800 56A/B etc. It could be as close as 15 years from now. Quantum cryptography needs to be ready well before then. Encrypted data could be copied and saved until there is a quantum computer to crack the encryption. The transition must be soon enough that any potentially compromised data is no longer sensitive when that compromise occurs.

Potential replacements for classical computers have pros and cons. Efficiency of these systems is comparable to quantum systems. Key sizes would be much, much larger than the key sizes used today. Some protocols would need to be changed due to upper limits on key sizes. The field has been getting more attention in the last few years. NIST has been working on the Post Quantum Cryptography (PQC) project. We are currently researching and

8 https://www.bbvaopenmind.com/en/quantum-computing/?utm_source=facebook&utm_medium=techreview&utm_campaign=MITcompany&utm_content=QuantumComputing
9 http://nvlpubs.nist.gov/nistpubs/ir/2016/NIST.IR.8105.pdf

presenting to interested groups. NIST has been presenting a paper on algorithms and interest is growing. We are looking for cryptographic algorithms that are believed to be quantum resistant.

The University of Maryland is actively involved in research in this area. The proposed timeline for research on algorithms calls for a November 2017 deadline for submission of proposals, and a three-year analysis phase. NIST will report its findings. NIST is looking for algorithms that are not vulnerable to quantum attacks. They are looking to find several good choices for new public key crypto standards. It is similar to past competitions, but there are differences too. Quantum cryptography is much more complicated. Submissions will include implementations. Evaluation criteria will be detailed in a formal call. Draft criteria will be open for public comment. Public evaluation of criteria will be strongly encouraged. Will also look at prior history analysis. Target security levels. The unit of work for a quantum computer has not been defined.

The Internet Engineering Task Force (IETF) has interest in the subject but has done nothing yet. The evaluation will be 3-5 years long in total. There will be two evaluation phases to consider all possibilities. Hash based signatures could be standardized in the next year or two. We are mainly focused on signatures and key establishment. We expect a transition in about 10 years.

Dr. Chen will continue to report to the Board.

Board Review

Earlier actions from today –

- The Board will want to consider how it can be of assistance to Mr. Groman's effort. (privacy council)
- **Next Meeting**: Have Peiter Zatko come to the next ISPAB meeting to discuss what his company is doing.
- **Board action:** Letter to the Chair and Co-chair of the Cybersecurity Commission also possible regarding recommendations. Chris will try to determine process for a letter. Letter in the next 30 days. Possibly an open call also.

Delayed voting on the following until Friday. There was no quorum present at the end of the meeting.

- Approving Meeting Minutes for March 2016
- Planning for Meeting in October 2016

Meeting Recessed

The Board adjourned at 3:57 p.m.

Thursday, June 16, 2016

Blockchain Protocol and the Emerging Ecosystem [10]

Dr. Ed Felten, Deputy U.S. Chief Technology Officer, Office of Science and Technology Policy (OSTP), The White House
Travis Hall, Ph.D., Telecommunications Policy Analyst, Office of Policy Analysis and Development, Telecommunications and Information Administration (NTIA), U.S. Department of Commerce
Andrew Regenscheid, Computer Scientist, Computer Security Division (CSD), ITL, NIST

Dr. Ed Felten, Deputy U.S. Chief Technology Officer, Office of Science and Technology Policy (OSTP), The White House

What is a blockchain? Blockchain is a distributed database that maintains a continuously growing list of data records that have been hardened against tampering or revision. It is made up of blocks where each block contains a subsequence of events. Each block has a unique predecessor or pointer. Blocks are made periodically (e.g., every 10 minutes on average in Bitcoin). The history of blocks grows over time by the addition of new blocks on the right end of the chain. Each block contains a timestamp and information linking it to a previous block.

Blockchain provides a public ledger that is tamper-evident. Any attempt to re-write previous history will be detected. If it is true that previously published blocks remain available forever, then the ledger becomes tamper proof. It doesn't solve all problems, only those where a public ledger helps. It is not very efficient if the user only seeks to publish information.

Who publishes the blocks? There are different models. In a central authority model, the authority creates the blocks, digitally signs and publishes them. In a quorum model, there is a well-known list of authorities who agree behind the scenes on the contents of the next block. They sign the agreed-upon block and publish it. Users trust any block signed by a quorum of authorities. In a fully decentralized model, everyone cooperates to agree on the next block.

Required properties of the ledger include "liveness", meaning at least one new block is published in every time increment; and "consensus," meaning no block has more than one successor (no forks). It is critical to avoid forks. If a central authority creates a fork, the authority is punished.

Within a quorum system, a fork is evidence some part of the authority misbehaved and must be punished. In a fully decentralized approach, it becomes very complicated. All participants must agree on the next block. However, there are issues which complicate the process: large numbers of participants where not all are known, and where some might be sock puppets for others; participants have incentive to manipulate the outcome; honest participation may be costly; and there is no assumption a single legal system will reach all participants. It may seem impossible in these circumstances.

In these situations, designing an incentive mechanism, so that if a large majority rationally maximizes their revenue, and assumes everyone else rationally maximizes their revenue, then consensus will result. The mechanism must be self-enforcing in the sense it relies only on cryptographic math and the laws of the universe, not on legal enforcement.

---

10 http://www.cio.com/article/3055847/security/what-is-blockchain-and-how-does-it-work.html

There needs to be a way to make this approach work. There is a solution in Bitcoin. It created a blockchain and digital currency together, creating a symbiotic relationship between the two. The blockchain gets used to record transfers of the digital currency. The currency gets used to pay people to make new blocks for the blockchain. The key idea is the payment is contingent on the new block being accepted as valid.

This seems to bring about consensus because the knowledge that if a block is created that forks the chain, it might not be accepted by others. Blocks that don't create forks have greater probabilities of being accepted. The incentive for users then, is not to create blocks that potentially cause forks in the chain.

Blockchains must either use a centralized or quorum approach or be used symbiotically with a digital currency. A "bare" blockchain that is fully decentralized is not something we know how to build.

What makes a valid transaction? The rules in bitcoin define a payment as – a payment of some coin from party A to party B if A owns the coin, and A has a digitally signed the transaction. Creation of new coins is valid if no more than 25 new coins are created, and there is only one-coin creation transaction in the current block. Transactions are accepted as valid if the greater majority believes the transaction will be accepted as valid. A transaction validity rule will be enforced if the greater majority believes it will be enforced. These rules can change whenever the greater majority believes the rules have changed. The reward for creating the block is made within the transaction. Blocks with transactions that are considered invalid will be ignored. The result is an incentive to avoid creating blocks with invalid transactions.

Coins come into existence because of the transaction. There is no backing and no inherent value. The value exists because of the belief that it will be exchanged for value later or demand. Bitcoin currency and blockchain use pure pseudonyms. A pseudonym is a randomly generated crypto key. Anyone can make a new pseudonym at any time, and can make as many as they like, whenever it's needed. It's easy to "follow the money" because all transactions are published on the blockchain. The identity of the actors isn't shown, only pseudonyms. Can pseudonyms be linked to identity? Sometimes, but it is complicated.

Soon, there will be fully anonymous crypto currency. Researchers have discovered ways to allow fully anonymous cryptocurrency. Blockchain consists of crypto blobs that no one can read. But an owner can prove to any third party they own a coin, and can transfer coins to another party. Blockchain itself conveys nothing of use to an analyst. A start-up company is building this now, called "Zerocash", or "Zcash". This will pose interesting policy challenges.

Travis Hall, Ph.D., Telecommunications Policy Analyst, Office of Policy Analysis and Development,

There is no central authority and significantly less control of the flow across borders and how it's used. Crypto currency as cash with metadata vs. anonymous cryptocurrency can be very significant. It is not traceable and can have some strange uses. Uses such as health care records management, real estate title management, IoT device management have been talked about but not used yet.

Public ledgers are being used, bitcoin could come in to help or assist with that. We don't know yet how it's going to be used. We also don't know the policy risks, yet. It's a bit early. The actual areas of concerns are not known.

Mr. Boyer: The biggest challenge, whether it's a social or technical issue is that governments can't control this currency.

The larger question is, can we trust the robustness, with the risk of tampering? These are probably social risks. Is it the right tool for the task of public key infrastructure (PKI) replacement? Can we keep it working? There are

technical issues around advanced applications. Scalability is also a factor. What about verification at this scale? If there are 10 thousand participants to one block, what about transaction counts?

It would be time consuming and painful; the download and verification takes so long. There is no really good answer to crypto currency. Downloading and verifying each N node is not sustainable with a very large number of participants. Larger groups of folks must then rely on smaller groups of folks to do all the work to do transactions without the middle man or intermediaries. In practice there are checkpoints that are universally accepted.

What is the global benefit? There certainly are nefarious uses. There are conversations to be had.

What's the intention of the government? Is it to trade dollars for bitcoins, like a foreign currency exchange? From a privacy perspective, with the use of pseudonyms, every transaction is visible, and visibly available.

Andrew Regenscheid, Computer Scientist, Computer Security Division (CSD), ITL, NIST

NIST is ready to address standards and guidelines to support Blockchain. It is able to help other agencies that have questions, regulations, standards, subject matter experts for cryptography. We need to separate hype from reality. It impacts financial, supply chain, health care. It uses fairly basic hashes, digital signatures, but is starting to change toward newer systems and better privacy.

NIST and Software Assurance

Presentation on a Secure Toolchain Competition [11]

M. Lee Badger, Group Manager, Computer Security Division, ITL, NIST (*Presentation provided*) present

Mr. Badger is presenting on a series of competitions to incentive development of tool chains for the construction of new software. The hope is better toolchains may avoid or suppress many programming errors that result in vulnerabilities. The idea originated in 2010 at a National Science Foundation (NSF) workshop. The idea existed for a while prior to 2010, when the opportunity arose to bring it into existence.

Each competition will pose a problem that the competitors must solve. The difficulty level of the problem needed to be determined so that progress could be measured. The difficulty level of the problems posed to the competitors increases over each test. The goal is to end up with reproducible results, potential technology improvements, and public data. If none of the solutions pass the minimum standards (having zero flaws), then there is no winner. If there is no winner, another competition is held with a different problem, but at the same difficulty level as the previous problem. The existence of flaws is determined using a mechanized test suite.

It is an objective test of programming skill. The competition is not interested in measuring people, but tools, and the ability of tools to help people to avoid making mistakes. People are flaw generators. For each competition where there is an actual winner, the difficulty increases for the next problem.

We want to improve the quality of software to reduce flaws. Every year, many flaws get reported to the National Vulnerability Database. Every year, the competitive scan report shows, there are thousands of flaws in software.

11 http://csrc.nist.gov/news_events/cif_2015/research/day1_research_1100-1150.pdf

There are databases on exploits and they're not always going only up. Sometimes, there are dips. It may represent the fact that these are social processes. Everything that happens may not get reported. The same level of focus may not be applied to these attacks all the time. There is always more malware being reported. It is a question of uniqueness. It is hard to say if everything reported is truly unique.

According to the Bureau of Labor Statistics, there are over three million people in the United States that can be classified as developers. We are trying to determine a place of good leverage. It is true that developers should be trained to write good code and make fewer mistakes. There are projects that focus on doing that. In addition, we also need good tools to help them. Mr. Badger's goal is to assist with the effort to provide those tools. There are all kinds of wizards and tools that help us with digital interfaces and other things, but not many assist with security.

What is a tool chain? It is a collection of software or hardware mechanisms that a developer may use to produce a software entity that executes on a particular platform. There are many platforms. How do you improve a software platform? If one wants to improve quality of software platforms, how can that be accomplished? There are millions of people involved, millions of devices, and many complex platforms. We believe incentivizing tools is a way to get progress in this area.

The competition cycle is about six months. It begins with an announcement. We want to have 4-person teams. This is because with a team of four people there is plenty of interaction to develop a solution to the problem. It facilitates planning and assignment of tasks. They will work on a non-trivial program in a short period of time. It is designed to take a team effort to produce a solution.

On game day, the problem to be solved is announced and the timer starts. The teams then work, and the program assignment must be checked in to the database by the deadline. Initially, we hoped the programming assignments would be these beautifully clear documents with solutions. It turned out it takes longer than that to describe what a program should do for the solution to be specific enough to be scored. The next day is scoring day. The programs are loaded into the programs to be evaluated for flaws. The submissions are scored, and if there is a winner, there is an award. If there is no winner, there will be no award. If there is a winner, the contest is evaluated for lessons learned. NIST would be able to codify what worked and what didn't work in terms of techniques sued to build the software and push that information out to the industry.

The goal is to identify and measure the most effective types of development tools. To do this, people should be allowed to bring any type of tool to the competition. The environment should not be restrictive. There is no influence from the sponsor. Scoring is mechanical. After the competition cycle, the scoring structure is published for public examination. It validates the scoring and allows people to learn from what was done. The results should be reproducible, and will be published so that participants and outsiders can use parts of the code and run them as they wish. It provides proof the problem was solved correctly, or perhaps problems may be discovered after the fact.

The challenge problem has three parts: functional specification, required security policy, and a problem-specific test suite. The test suite is written with knowledge of what the problem is. It attempts to expose what might be wrong if one of these test suites were implemented.

Originally, the thought was the challenge problem could be described very briefly. It took twenty pages to describe a problem. We are working on ways to describe challenge problems more concisely. There are three types of software for challenge problem types: Command line interface, mobile, and web (html based applications). Each type of problem has particular requirements for the solution. The participant provides a deployable virtual machine image. The test infrastructure boots the machine, allows it to stabilize, and then it runs the tests.

The mobile challenge problem involves developing a news app. It must meet the criteria displayed in the problem. The teams provide the app. The competition provides the server the app interacts with. The user interface must meet stated criteria: The solution must be an Android application launched from the Android home screen. It must be able to receive files, interact with the network and accept user interface input. It should also perform certain functions, generate any required data or protocols, and it must be in Android package format.

There are 20 pass-fail functional tests, 20 pass-fail security tests and additional fuzz testing. A faster submission time serves as a tie breaker in the event of a tie in scoring. Completion time under 10 hours, wins. Currently, we are working toward developing a scale to judge software complexity. We built and tested a testing architecture using a dry run. The dry run consisted of eight tests, with a total of 12 developers. It turned out that submitting virtual images together or even serially, was too heavy a load for the system doing the automated testing.

The testing produced a series of lessons learned. Teams should be warmed up (know each other, and have practiced together), have sufficient context for the problem, etc. It is difficult to get a good idea about tool efficiency if the team has not really formed. It became important to provide more context, so that when competitions are announced, they provide information about what skills are needed.

They have been in the preparation phase. They have formulated, documented and implemented the eight problems. There has been a dry run competition held at NIST, with the help of the National Cybersecurity Center of Excellence (NCCoE). During the testing they learned the testing infrastructure needed to be simplified. They are working with the contractor to arrive at the proper documentation. They provide weekly statuses. There will be a re-design phase following finalization of the documentation. Following their own inspection, they discovered the design could be clearer. We are trying to change the world with better tools, and fewer security incidents. Eight apps have been implemented for competitions. Three Android apps have been implemented as a result of the competitions. We are hoping to lift the state of practice.

Mr. Boyer: These concepts are extremely interesting. How do you move through the software and validate that it's legitimate?

Mr. Badger: I believe there are commercial applications that handle this. In the last ISPAB meeting, we talked about open source code and vulnerabilities, and being sure it's not some sort of malware.

Mr. Badger: The competition is not about searching for malware. It is crucial to do, but it is about helping people with good intentions, and not creating flaws and vulnerabilities. We let people use any library or framework they want to use. There are very powerful, possibly malware ridden component to get the assignment done quickly, they can do that. If the malware causes one of the tests to fail, then it could be said the malware was detected. But it's not what the tests are about.

Mr. Boyer: I was thinking from the standpoint of, if the code was more secure from the beginning than whether someone found a new way to exploit it.

Mr. Badger: Yes, that is the case. We are hoping to lift the state of practice.

Mr. Greene: Do you see this filtering back into the educational environment?

Mr. Badger: Yes, that is the goal. It involves how techniques are used. Competitors must allow their software to be viewed while it is in use.

Presentation on Foundations of Software Assurance [12]

Paul E. Black, Ph.D. Computer Scientist, Software Quality Group, Software and Systems Division, ITL, NIST

There are a lot of software projects at NIST. Today, we will be discussing two projects: One is called the Software Assurance Reference Dataset (SARD), and the other is the Bugs Framework (BF). SARD is a public repository for software assurance cases with known vulnerabilities. We have collected over 140,000 cases in multiple languages. We have received contributions from many people to put in the slide. For background, here is why it's useful.

Programmers write static analyzers to search for vulnerabilities in source code. What we want to know is, does this work? Do these static analyzers find bugs in programs, and how do we know? In order to determine an answer to that question, they test with programs with known bugs. They run the static analyzer and see what the result is, and determine what was found, and what might have been missed. There is a preference to production code for testing. Ideally, there should be a statistically significant number of known bugs for the software to find. That way percentages can be determined. Unfortunately, they don't get anything close to that. There are three different types of test cases: software where not all errors are not known; production code with known bugs (has vulnerabilities tracked back to particular lines of code); many cases with known bugs. It is very labor intensive to trace vulnerabilities back to lines of source code. There are only several hundreds of those in the SARD repository at present. There are also cases with known bugs.

SARD has contributions from many static analysis researchers. NSA and Intelligence Advanced Research Projects Activity (IARPA) created and gave thousands of test cases from their work to SARD in order to distribute to the world. They are adding over a thousand test cases from Toyota. The test cases are to be made available to everyone. Additional content came from MIT in the form of static analysis toolmakers. It has been helpful for tool qualifications. When analyzing large volumes of software, it can save significant amounts of time. There is much more work that happens with SARD, but there is not time here to discuss.

The Bugs Framework (BF) is a descriptive language Mr. Black's group is trying to develop a language that is intended to precisely describe bugs in software. Current bug descriptions exist, but they have significant problems. They tend to be imprecise and inconsistent and lacking detail. Some combinations are listed and others are. It does not distinguish between types of bugs. Each bug class has a cause, attributes of the bug, and what are the consequences. These are represented as directed graphs. They are pushing for a level of formality in the framework so that when definitions are derived, we know with certainty what techniques and tools will prevent particular occurrences and it can be proved that they work. There are four classes in the framework so far: injection, control of interaction frequency (log in limits, etc.), information exposure, and buffer overflow.

As an example, there are certain questions to be asked when a buffer overflow occurs: Was it caused by a read or write activity; was it out at the high end or low end of the buffer; is the array in a stack or heap; how far is the array outside the boundary; etc. Attributes of flaws display in the framework so that remedies can be determined. There are two errors that cause buffer overflow: Data exceeds the array, or the program points to the wrong location. The consequences from these errors are numerous. It took some effort to describe what was going on from other people's descriptions, and be able to write about those in a consistent way.

12 https://samate.nist.gov/SARD/ https://samate.nist.gov/BF/

What is the Bugs Framework good for? It precisely explains why techniques work in some cases and not others. It describes vulnerabilities more clearly. It helps programmers to write better code. It accurately states the classes of bugs that software assurance tools cover, or do not cover. It makes precise explanations possible. Heartbleed was caused by data exceeding the array because the input was not checked properly. We are looking to enable secure tool chains. We hope it will be a stepping stone for other efforts.

Is there any uptake in people using this nomenclature? It is not being used yet. There is possibly a new project that will use the Bugs Framework in the near future. It is somewhat less useful in a stack because there is less randomization.

Presentation on Combinatorial Methods in Software Testing [13] [14]

Rick Kuhn, Computer Scientist, Computer Security Division, ITL, NIST

Mr. Kuhn will be discussing methods and tools to make software testing more efficient and effective. There is a great deal of industry interest in this work, and many companies are using this method.

The primary application of these methods is for functionality testing and errors and security vulnerabilities. It is true that two thirds of all vulnerabilities come from software faults. Buffer overflow is a classic example of this type of vulnerability. It can be used for modeling and simulation as well as for performance tuning. This is very effective for systems with a large number of components that interact.

A couple of years ago, we put together a cooperative research and development effort with Lockheed Martin. As far as we know, it is the largest publically documented trial use of combinatorial methods. The results were really successful with a 20 percent savings in test costs, and an average increase in test coverage ranging from twenty to fifty percent. It was measured across eight projects in database and avionics. Twenty percent represents a solid result.

More specific applications show an even greater improvement. Rockwell Collins applied NIST Methods and tools on testing to FAA life-critical software on aircraft. They found very large cost reductions when these methods were applied. It is important because typically, software development is about half the cost. In civil aviation it is approximately a 7:1 ratio. It was a very big improvement.

How did we get here? What is the empirical basis for what we are doing? We started looking at causes of failure in software in 1999. Many areas were looked at such as logic errors and calculation errors. One of the things we looked at was interaction faults in medical device recalls. It is a two-way interaction, consisting of testing all pairs of values to find this fault. If both pairs were found together, they would find the fault. If factors were not found together, the fault might not be found. There could be more complex inter-actions. There can be three-way interactions, and three-combinations. No one had looked at beyond 2-way interactions before.

13 Combinatorial testing, which derives from the field of Design of Experiments, has attracted attention as a means of providing strong assurance at reduced cost. Combinatorial testing takes advantage of the interaction rule, which is based on analysis of thousands of software failures: most failures are induced by single factor faults or by the interaction of two factors, with progressively fewer failures induced by interactions between three or more factors. Thus when faults are triggered by a combination of t or fewer parameters, testing all t-way combinations of parameter values can be considered pseudo-exhaustive, and can provide a high rate of fault detection. The talk explains the rationale, method, and tools available for combinatorial testing.

14 http://csrc.nist.gov/publications/nistbul/itlbul2016_05.pdf

What this tells us is, about 65% of faults are single factor faults. More than 95% of the total were caused by one or two factor faults. This represents one industry and one application. We then looked at open source Apache server, which had much more complex faults than the medical device software. There were fewer single factor faults than in the medical device software. It is also possibly related to the number of users. The server software has many more users than the medical device software. We also looked at browsers and found similar curve with only up to six-way interactions.

We looked at NASA's distributed database and found that the curve was similar to the medical devices, but slightly better. That was the first time it was tested. Curves for SQL, and TCP/IP stacks were very similar. No matter how many we look at, we always find a certain curve. No one has ever seen a failure where greater than six factors are involved. Most factors are caused by either one or two factors. This is known as the interaction rule, which says that when all faults are caused by *t* or fewer variables, then testing all *t*-way combinations is pseudo-exhaustive and can provide strong assurance. It is almost always impossible to exhaustively test all possible combinations. The interaction rule says it is not necessary. On a practical level, it works very well, and the data supports it.

For testing all possible combinations, pseudo-exhaustive testing is very effective. In computing the number of tests for a combination, the combinations can be arrived at mathematically. The bad news is tests increase exponentially with the number of variables we are interested in. The number of test combinations covered grows very quickly. On average, a very high percentage of failures can be detected with three-way coverage. Untested combinations are where problems will occur. We won't know what that code will do. Measurable values relate directly to assurance of software. Testing thoroughness can then be defended.

How is this used for testing assurance? It is useful for providing a measurable value with direct relevance to assurance. When it's asked how thorough a test set is, a defensible answer can be given. It provides a measure of the input space that was covered. It also provides a means for contract monitoring and determining quantitatively if testing has been sufficient.

Oracle-free testing provides a way to do testing without conventional test Oracle. There are ways to accomplish this currently, fuzz testing is one of those ways. A new method is to use two-layer covering arrays. It only requires definition of equivalence classes. Once the equivalence classes are defined, everything else is automatic.

Dr. Romine: In general, the way you describe this, the software generates close to an optimal cover of the test bed.

Mr. Kuhn: We can generate a very small number of tests. There is no known way to determine what's optimal.

Dr. Romine: Given the set of tests that have already been run, is there a way to maximize the utility of the next test, and generate next tests that expand the coverage the most, and is the software capable of doing that?

Mr. Kuhn: Yes, the software does that. We can take an existing set of tests, and expand it to any level of coverage we want. One tool will automatically adjust the covering array, the other will expand the level of coverage.

High Performance Computing (HPC) Security

M. Lee Badger, Group Manager, Computer Security Division, ITL, NIST

Tim Polk, Assistant Director, Cybersecurity, National Security & International Affairs Division, Office of Science and Technology Policy (OSTP)

Mr. Polk is a NIST detailee who has been working on cybersecurity and the National Strategic Computing Initiative (NSCI). This is the first time the NSCI has been reported to ISPAB. It is the right time to talk about high performance computing security. The NCSI was created by an Executive Order - National Strategic Computing Initiative in July 2015. Two years of work were done prior to the executive order being issued with other agencies being involved. All agencies have investments in high performance computing.

There a number of nudges in the correct direction that had been received over the previous years. The OSTP pointed to the President's Council of Advisors on Science and Technology (PCAST) is a (Federal Advisory Committee Act) FACA run by Mr. Polk's office. Through their bi-annual reports, they have been pushing for a national strategy on supercomputing. It was one of the drivers that led to the present time. There was some difficulty in selling and really marketing what they were seeking to achieve.

A whole-government approach was needed to explain what was needed to explain why it was important not just for agencies, but for the whole nation.

The NSCI is intended to be whole government and whole nation. It is a strategic program that is not about building new platforms, and not any specific technology. It is intended to augment economic competitiveness, and staying at the front and have the most advanced computing capabilities. This initiative is not intended to be business as usual. There are a lot of agencies investing a lot of money and resources. It is intended to accomplish more than merely meeting mission requirements.

There are five strategic objectives:

1. Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of the current 10 petaflop systems across a range of applications;
2. Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing;
3. Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached. (The post Moore's law era);
4. Increasing the capability and capacity on an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development;
5. Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government, and industrial and academic partners.

Key themes derived from strategic planning:

Strive for convergence of numerically intensive and data intensive computing. Keep the U.S. at the forefront of high performance computing capabilities. Stream line high performance computing development. Make it readily usable and accessible.

Related initiatives: materials genome initiative, advanced manufacturing initiatives, the Brain Research through Advancing Innovative Neurotechnologies (BRAIN) initiative, precision medicine initiative, National big data R&D, and the National Photonics initiative.

What would success look like? It mirrors above.

Security was not included in the executive order. It was raised at the first executive council meeting. It was established as a cross-cutting sixth objective. We cannot succeed in the other performance objectives without security.

M. Lee Badger, Group Manager, Computer Security Division, ITL, NIST present

NIST will host an HPC security workshop in September in the hope to form a community for HPC. Many questions exist regarding HPC. There was an organizing meeting in May. It is now in late September. A community of security experts and HPC experts need to be built. Agenda topics: current use cases and practices, review the framework, and discuss threats followed by breakout sessions and bringing groups back together to report on findings. This is the first of what is hoped will be a series of meetings.

Understanding what can be done with HPC for the enterprise, the data available is not being exploited by those enterprises that have HPC. The hope now is to add the security at the beginning instead of adding it later.

Legislative Updates Relating to Security and Privacy

Jessica Wilkerson, Professional Staff Member, U.S. House of Representatives, The Energy and Commerce Committee

Ms. Wilkerson covers cybersecurity issues for the House Energy and Commerce Committee. There are seven topics to discuss today in the area of vulnerabilities: Are there things that can be done, etc?

The first area is, internet cybersecurity concerns. We are working to ensure first awareness. There have been intergovernmental meetings with agencies on cybersecurity issues. It is a big deal for various networks.

In June 2015, there were two incidents following a Lenovo issue with self-signing certificates that caused concerns. It involved issues with the trust chain. A tool was issued to resolve the problem. NTS Holding accidently started issuing certificates for sites they had no authorities for. It was fixed in a couple hours. They wrote to the four big browsers about named Government certificates as well.

The cryptographic round table was held in October 2015. What are the paths for removing outdated cryptography from the system? Smartlock and cryptography may not be good in five to ten years. Cryptographic issues continue, and needs to be updated as the internet of things continues to expand. We want to start having conversations now, rather than ten years from now. Educating people about cryptography is not a one-and-done affair.

Automobile and medical device meetings were held in February and March 2016. A lot of action is happening in these areas. We will see more about efforts in these areas as time goes on.

A report on Health and Human Services (HHS) cybersecurity was released in August, 2015. Security and operations clashed, with operations winning out. A rash of security incidents were related to that situation. The recommendation from the report was to elevate security. The HHS CISO is solely responsible for security. The Bill is HR5068. Hearings were held to discuss the bill in May. In April 2016, a bipartisan bill was introduced to make the HHS CIO a peer to security in the agency organization chart. There will be a report to Congress on why they need this a year later.

The Encryption working group is in an early stage of information gathering. They are working with academia, Federal Government, state, industry, and law enforcement. We want to move past the law enforcement vs. technology perception, and move past the one-size-fits-all solution. A combination of solutions to suit federal, state and local law enforcement needs must be developed. Concerns for each group will be different.

Automobile cybersecurity is a major area of interest for the committee. The Automotive Information Sharing and Analysis Center (auto-ISAC) is not active at this time. The committee is investigating what role the auto ISAC may be able to take. There are no answers at the present time.

Mr. Boyer: ISPAB will continue its interest in encryption, and we hope to revisit this topic in the future.

Updates on NIST Cybersecurity Framework [15]

Matthew Barrett, Program Manager, Cybersecurity Framework, NIST

The last time Mr. Barrett spoke before the Board was in September, 2015, and provides updates to ISPAB about every six months. Currently, he is involved in a great number of free market initiatives including: working with the Coast Guard on to define a cybersecurity framework; he is also involved in the Health Insurance Portability and Accountability Act (HIPAA) security framework; and, working with the tax system and application of the framework controls to prevent identity theft, and working with other laboratories on discreet manufacturing. There are a number of initiatives within the framework in process now. Outreach efforts continue with international, small and medium businesses. There are 28 countries in the dialog on the framework. Italy adopted the entire framework, and also extended a few areas.

Small and medium business has training efforts in process such as assisting organizations in understanding why cybersecurity is important, and how to manage cyber risk, and approaches to readiness to apply the framework. There was a training effort held in South Dakota for communications companies.

They are continuing with regulatory dialogs; the Cybersecurity Enhancement Act is central to those discussions. They have authored a request for information on the framework, its use, etc. That led into the question, Is it time for an update? What is the long-term maintenance going to look like? The RFI was published in December, 2015. The analysis was published in March, 2016. All of that was preparatory to the workshop held on May 27th, which

15 https://www.nist.gov/cyberframework

was held to get everyone together to talk in person. The framework group is busy with requests for information on framework use, how it's being used, etc.

There are future plans to publish guidance on the framework approach, NIST's role, how to make changes, etc. We also intend to publish self-assessment criteria. Key concepts come from the Baldridge Excellence Builder. The new iteration of the framework and the self-assessment criteria will be published together. It will be version framework 1.1. A framework for supply chain and others will be included. There is a change control log to track possible changes for future editions of the framework. All suggestions were reviewed and a short list selected. The update is projected for early 2017.

There is an interesting dynamic to Baldridge awards, and whether companies become targets on recognition given by Baldridge. The Baldridge award is more a recognition of excellence, rather than an award in the usual sense. People understand the value of a recognition program. It seemed an easier place to start. We will look to see how things develop. We are evolving the framework at the same time the Baldridge recognition program is going on. A minor update to the framework will be published in the early part of 2017.

How can the framework be used for challenges identified by the commission to enhance cybersecurity? How can the framework be applied in different environments? Can companies doing basic IoT services use the framework to assist with cybersecurity? We are working with a short list of items to determine which will make the cut for the update. Common themes in feedback would be interesting to see. Can we determine what is uniform and what is marginal?

NIST Updates

Matt Scholl, Chief, Computer Security Division, ITL, NIST
Kevin Stine, Chief, Applied Cybersecurity Division, ITL, NIST

NIST had actions to address under the CNAP. They are publishing a white paper on best practices for PIV card use. Many back-end systems do not support PIV. NIST has discussed capabilities for establishing middleware. One of the other deliverables was to provide more guidance on the recovery aspect of the NIST framework. There was guidance on set-up, and initialization but not on recovery. The draft was sent out for comment a couple weeks ago. They are also looking at incentivizing organizations in the use of the framework.

The rest of the work in CNAP is supporting other agencies. It also looked at confidentiality breaches, and reacting to and recovering from the use of ransomware. There is a list of documents that have come out since April. The cryptography staff is expanding at the NCCoE. They continue to look at deeper and foundational issues and research areas in security and privacy areas. They also will be examining security and privacy in social media and big data. There are cross-disciplinary groups working to look at security.

Mr. Stine is working to build a division leadership team. That process has been going relatively smoothly. They have made some new hires in cryptography, and added guest researchers. Mr. Scholl updated the group on the framework. One of the opportunities with the framework is where or how it intersects with other NIST work that is going on. The NICE is ongoing, and is about at the halfway point. Regional Alliances and Multi-stakeholder (RAMPS) to stimulate cybersecurity closes July 12th. There are regional needs from a cybersecurity workforce perspective. All facets of industry can get together to work on the needs of the cybersecurity effort in local areas.

There will be some new there will be some new National Strategy for Trusted Identities in Cyberspace (NSTIC) pilots in September, but awards will mainly be made at the end of the fiscal year. Application review is going on now. We will be releasing the first version of the identity ecosystem framework. The registry was announced, and it allows for organizational self-assessment and self-attesting the framework is being used. A public draft of SP 800-63 has also been released. Feedback has been positive and increasing. We are incorporating more privacy and usability requirements over the course of the process. Privacy efforts go on outside of cybersecurity efforts.

We anticipate the releasing a special publication later this year, and it is expected to contain a companion roadmap, more privacy controls. Some federal agencies have provided feedback. Most has been positive. They are working with the Federal Privacy Council. Not all privacy work is happening in the Applied Cyber Security Division. A number of NIST privacy team members are working with the privacy council.

The NCCoE is continuing to gain momentum with a retail-focused workshop held at the University of Alabama; in order to help retailers, implement stronger protections. Topics included authentication for e-commerce; stronger authentication for ecommerce transactions. There is work on a practice guide for secure email.

Health care work continues. Kevin Fu drove a lot of interest in health care and health care devices. On the partnership front, the last several months have been active as a result of partnerships. NIST is now partnering with VMWare, and discussions with the communications sector were very positive in working with them at the NCCOE.

Cybersecurity workshops are planned for the end of the fiscal year, including use of blockchain for healthcare and healthcare records. We were approached by HHS and are working with them on blockchain and cryptography. HHS is looking for research topics.

Privacy and Civil Liberties Oversight Board (PCLOB) Updates

David Medine, Chairman, PCLOB

Mr. Medine Will be with PCLOB until July 1st. The Board is working on Executive Order 12333. They are doing deep-dives on three activities – 2 CIA, 1 NSA. These activities are classified. The report will be out by the end of the year.

The Board is engaging in other activities, including Section 803 reports. They are working with various agencies to see if the reports can be made more informative. Agencies are reporting to the Board and to the public. Major activities include implementation of recommendations from the report. Mr. Medine is leaving the Board July 1, leaving four remaining members. Congress gave the Board a dozen new hires. Hiring is on track to be completed before Mr. Medine departs. The issues are technological, legal, and policy-related.

Things will be more complex next year. Departing members will not be replaced this year. The next president will have a major say over the Board. Staff will be long term. There may be detailees, but most will be permanent hires.

There is language pending to limit the Board's oversight to U.S. persons. PCLOB currently covers intelligence of non-U.S. persons. If that stops, there will be a gap in oversight. Countries can exchange intelligence information. The House is going to vote on requiring a court order for gathering information of American persons. Section 702 is

up for a vote. There will be a public report by the end of this calendar year. It is a huge body of work and a big change in how these issues are addressed.

Mr. Boyer: Can the Board assist with these efforts?

Mr. Medine: The Board can make recommendations for filling the slots. The challenge will be for the next president to fill those positions.

Meeting Recessed

The meeting recessed at 4:30 p.m. The start time for Friday will be 8 a.m.

Friday, June 17, 2016

The Chair opened the meeting at 8:02 a.m.

Internet of Things – Part I

Nathaniel Beuse, Associate Administrator, Vehicle Safety Research, National Highway Traffic Safety Administration (NHTSA)
Karen Jagielski, Senior Attorney, Division of Privacy and Identity Protection, Federal Trade Commission
John Morris, Jr., Associate Administrator and Director of Internet Policy, Telecommunications and Information Administration, U.S. Department of Commerce [16] [17]
Travis Hall, NTIA

John Morris, Jr., Associate Administrator and Director of Internet Policy, Telecommunications and Information Administration, U.S. Department of Commerce

NTIA's primary work on the internet of things has started in the last few months, with a request for comments regarding what issues and concerns are raised by the internet of things, and what actions the government should take to address these issues and concerns. They received more than 130 substantive comments were received. There were particular issues that were raised, including privacy questions. How do we in the US respond to privacy? Does the internet of things raise the same sort of privacy questions as the internet in general? The current administration has voiced strong support for privacy legislation.

There are some IOT specific security issues. If a device is installed, and the device is later found to have a vulnerability, can it be patched, or is security a one-shot event. Do best practices exist for upgrade practices and disclosure to consumers on what can be done in these instances? Many stakeholders were urging examination of upgrade practices, and best practices regarding disclosures. We continue to examine these areas.

NIST and NTIA have ongoing technical research focusing on the internet of things. NTIA has a smart cities effort looking at radio frequencies, and number of transmitters in a certain area. The NTIA lab in Boulder, CO has a smart cities effort to look at and assess the impact on radio spectrum usage if we really want to move toward a robust internet of things. What is the impact of large numbers of transmitters in a small area? Generally, are there specific research topics that the government ought to engage in that it is not doing presently?

NTIA collects data on Americans use of technology and broadband access, etc. In the most recent data collection in July 2015, the NTIA contracted with the Census Bureau to add questions to the population survey. It is one of the largest surveys in the world and includes three thousand households. Only a small percentage of households currently have appliances online. The next data collection will be in 2019. We will be able to evaluate how fast things have grown. A green paper on IoT issues will be released in September or October of this year. It will seek to articulate what needs to be done rather than try to set policies.

---

16 https://www.ntia.doc.gov/files/ntia/publications/fr_rfc_iot_04062016.pdf; https://www.ntia.doc.gov/federal-register-notice/2016/rfc-potential-roles-government-fostering-advancement-internet-of-things; https://www.ntia.doc.gov/press-release/2016/us-department-commerce-seeks-comment-potential-policy-issues-related-internet-thi

17 https://www.ntia.doc.gov/federal-register-notice/2016/comments-potential-roles-government-fostering-advancement-internet-of-things

Ms. Levin: Depending on disclosures is not effective. The struggle to deal with these issues has been going on for some time. The FTC went the disclosure route with the new marketplace. It is good there is only a small percentage of households because it gives you more time to figure out how to deal with these issues. These same issues existed in the old electronic market place, and it was a struggle to deal with them then. How far do you think disclosures will go?

Ms. Jagielski: There are certainly ongoing issues on how to make consumers aware of what data is being collected and how it's being used, is a challenge. We believe we can come up with a reasonable way to give consumers the information they need. Commerce looks at consumer facing products. There has been a learning curve and the effort to keep up with change. We are educating ourselves while being proactive and encouraging others.

The effort of privacy is a constant evolution. In the next five or ten years, Congress will revisit. We know what needs to be done. It's not rocket science any more. The action is really in the agencies in the absence of Congressional action. The "wait and see" course of action is not useful any more. We've waited and seen, while the risks get greater and greater.

Ms. Jagielski: The FTC has been asking Congress for privacy legislation for over 20 years, and will continue to ask. We have asked for broad principles for technology and privacy, and the problem is that often it takes and event to create change. In the absence of legislation, there is then the need to make the law a bit elastic to get it to where it applies in the current world. Fortunately, the FTC Act is pretty broad. Without the sort of legislation that specifically creates authority in an area, privacy battles are fought and re-fought. The Wyndham Hotel case is an example. Responsible companies generally want to do the right thing. The FTC works with industry.

Mr Robach: Is there not a tension between wanting to provide security, and what the lawyers want to say? We want "reasonable security". It becomes a means to evaluate the entire system a company may have in place. There will be breaches, they are not avoidable. Good players then begin to set what is reasonable. Industry is in the best position to figure it out. "Reasonable security" applies to service providers and hardware and software manufacturers. There are specific exclusions from the FTC Act such as banks and a few others. Common carriers are included in the authority of the act.

Nathaniel Beuse, Associate Administrator, Vehicle Safety Research, NHTSA

There are two types of IoT communications with vehicles: Vehicle-to-Vehicle, and Vehicle-to-Infrastructure, along with everything else connected to the car by consumers or manufacturers. They are looking to build out a functioning prototype to be ready to go on day one. It has not been built as yet, and it's not needed yet. Research and testing going on in parallel slowly at this time. We are trying to discover vulnerabilities before going to production. NHTSA will be making a proposal on vehicle security this year. They are looking to start with supporting 17 million vehicles. Research will include searching for vulnerabilities, prior to the final version of the proposal.

The basis of the system is PKI. It is heavily focused on privacy. There is some tension between security engineers and privacy advocates. Having less privacy would streamline security somewhat. The engineers are still looking to keep privacy as strong as possible. The processes are separate and neither knows what the other is doing. There is some education that still needs to happen.

It is next to impossible to chase down every entry point into the vehicle. From a security standpoint, testers are focusing on the vehicle itself. They are looking at the standard things, separation of functions, watchdog algorithms, blocking certain functions.

There was a forum in D.C. in January this year and learned many lessons from that. Looking to develop best practices for the auto industry from the information that was received. They spoke to the Food and Drug Administration (FDA), FAA and about medical devices. They tried to come up with a best practices document for automobiles.

There is a trust issue within the issue regarding sharing information. An ISAC was started in January 2016. It's running but the challenge is how to share information. There will be a learning curve, as with all ISACs. The hope is to have more people involved. We want to have the ISAC function, and function well. The ISAC has put out a best practice document. NHTSA has not examined it as yet. The ISAC practices will probably be more detailed than NHTSA's will be. Documentation will be key during this process. We are committed to doing more on privacy.

Oversight of all these pieces may fall to multiple agencies. Conversations of this type are going on with Google, Lyft, and others that may make automobile equipment. Apps on smartphones can create vulnerabilities on WIFI on vehicles. WHTS will not police peripherals but focus on the vehicle's response to those things.

There are a lot of security practices already out there. Identifying critical systems and making sure they are isolated from each other is crucial. May need to err on the side of doing more rather than less. In January, all 17 automakers signed the automotive safety agreement to make cybersecurity a priority. Each company has a different level of maturity, and responsibility for security within companies still needs to be determined in some cases.

Approaches from different agencies need to be compatible. Government and industry must communicate enough not to stifle innovation for consumers. Best practices can become too proscriptive. FTC and Gramm-Leach-Bliley arrived at the right level of guidance, so that everyone could operate. It may not be as hard as it seems.

There needs to be a means to be transparent. How can we be transparent when selling devices? There needs to be a way to talk about these things. FTC and Gramm-Leach-Bliley notices can offer examples for how to do disclosures. Key disclosures up front make it easier for consumers. There needs to be consistency in how disclosures are constructed.

Has there been exploration with insurance companies of security in vehicles, meaning if security upgrades are not applied, will that reflect in the driver's rates? It has been discussed. Insurance is still struggling with how to value cybersecurity.

Ms. Jagielski: A disconnect remains in the industry. Auto engineers are not security experts. Change is happening, but slowly. Cars can be looked at as one-ton or two-ton computers. The FTC site, www.identitytheft.gov has re-vamped its website. There are new forms and letters for consumers to use if their identity is stolen.

Mr. Hall: Disclosures are still being clarified. The issue is how to be transparent when devices are sold, and how can consumers be assured they know what they need to about a device they buy. They need to understand what to do if a bug develops.

Ms. Levin: The FTC is a good example of how disclosures should be done. Lawyers will need to be involved. Communication experts can assist with the language. Insurance companies can also be an example. Coverage can be impacted if certain components are not on the vehicle.

Mr. Hall: We are talking with insurance companies on this issue. In Europe, they are looking to come up with an industry standard everyone must meet. They are still struggling with how to value cybersecurity.

Ms. Jagielski: A disconnect exists across industries. There are a lot of automotive engineers that are not computer scientists. Consumers want cool features, and it was not a problem until 2010. A car is essentially a two-thousand-pound computer because it has so many lines of code.

The intent is to show what we can do before we try to present anything to the American people. We are hopeful that the worst case scenario is denial of service instead of someone getting into the system and messing with the brakes. The DOC and the FTC have provided comments and worked with someone who wrote the staff report. We are making sure that the conversation is happening.

The FTC is able to look at consumer-facing issues and the NHTSA is looking into the vehicles on American roads. We are looking into the conversations to see if the right topics being discussed to protect the consumers of devices. Does it make sense to break down the IoT, or to have a broader approach? In some respects, it does make sense; but with privacy there are no real cross cutting issues.

The approach is to look at conversations happening across the government and industry, and see where there are gaps, etc. The goal ultimately is to protect consumers. Does it make sense to silo conversations to a certain degree or have a broader approach? A broader approach makes sense. It involves continuing to have these types of conversations. The economy ultimately serves consumers.

Mr. Boyer: Can the Board offer thoughts?

Mr. Hall: The Board may respond to the RFC.

Presentation on U.S. Government Accountability Office (GAO) Reports

- *Information Security: DHS Needs to Enhance Capabilities, Improve Planning, and Support Greater Adoption of Its National Cybersecurity Protection System* [18]
- *Critical Infrastructure Protection: Sector-Specific Agencies Need to Better Measure Cybersecurity Progress* [19]
- *Federal Information Security: Agencies Need to Correct Weaknesses and Fully Implement Security Programs* [20]

Michael Gilmore, Assistant Director, Information Security Issues, U.S. Government Accountability Office (GAO)

The panel will be discussing three GAO security-related reports. They are distinct in the areas they cover. The first report is the review of the National Cybersecurity Protection System (NCPS). It has four system objectives: Intrusion detection, intrusion prevention, analytics, and information sharing. They wanted to understand the scope of implementation across agencies. They focused on five different agencies and each agency's use of the system.

In terms of intrusion detection, EINSTEIN is a signature based system, so it is limited in what it can do. NIST recommends a three pronged approach to the evaluation which includes signature based, anomaly based, and

18 http://www.gao.gov/products/GAO-16-294
19 http://www.gao.gov/products/GAO-16-79
20http://www.gao.gov/products/GAO-15-714

others. They looked at signatures to see if they matched common exposure vulnerabilities and others. They used multiple vendors. There was a mixed result. APT had better coverage overall. Understanding the purpose of the system is important in evaluations. It was not intended to pick up everything. Customers vary in terms of network security. Some agencies may need more or less review.

There is limited ability in Intrusion prevention. What EINSTEIN does is system blocking and some other functions. Its capability will expand. A variety of analytical tools are present. Information sharing is still to be developed, as requirements have not been approved as yet. An executive roadmap exists to lay out activities until 2018. At the time of the audit only five of the systems were using the intrusion prevention system. Internet service providers and companies were having challenges running classified systems in unclassified environments.

Why are agencies not sharing vulnerability information? It is a very fast moving area. Information presented here is already outdated. There is now an active automated information sharing system that includes government and industry. We are working with the first EINSTEIN 3a carrier to get them authorized to work in the environment. A proof of concept for a non-signature based capability is in process, to move into production next month. Also, there is a mandate for all agencies to adopt EINSTEIN 3a by the end of the calendar year. There is a lot of activity. EINSTEIN does not see vulnerabilities. Other programs provide the capability to detect vulnerabilities, and allow us to make more informed decisions regarding prevention. It creates a more robust picture.

The critical infrastructure protection report looked at nine sector-specific agencies that are involved with critical infrastructure. We are involved with the private sector to best understand how to improve cybersecurity. We wanted to determine how sector-specific agencies perform their job of determining risk. Treasury's work with the financial sector is a leading example. Only 11 of the original 15 agencies felt cybersecurity risk was important. Because of the timing of the report, they had to rely on the 2010 sector-specific plans as they were most current. The 2015 plans were supposed to come during work on this report. Some of those have yet to be published.

We also wanted to determine what actions had been taken. All had taken some steps, even those which did not deem security important. Most had analyzed dependencies and identified risks. Performance management is a difficult thing in the cyber security world. Fifteen sectors and 8 agencies were the focus of the report. Only three of those had taken some action on how to measure progress. The others had not taken any real action. How is progress on threats actually measured? It's difficult to do measurements as there are not many events to measure on. Metrics have been a major challenge. We can only demonstrate work with agencies.

The FBI's Use of Facial Recognition Technology

Diana Maurer, Director, Justice and Law Enforcement Issues, GAO

Ms. Maurer has worked in a variety of areas in national security. A report released this week stating the FBI needs to do a better job with privacy in its use of facial recognition technology. The FBI has not put the proper value on ensuring information is accurate, or ensure the privacy of individuals is protected.

The report looked at three issues: First, how does the FBI use facial recognition technology? There is a lot of information out there in the public domain that is partially correct. Part of the issue is rooted in the lack of transparency about what the information is being used for. There are two general ways the FBI uses facial recognition technology. One is the next generation identification interstate photo system. It involves an evolution of the way the FBI uses fingerprints. Originally fingerprints were printed on paper and stored. In the 1990s,

computer technology came on the scene. Fingerprints are now digital, and now include images (mainly mugshots). The FBI has collected around thirty million images over the course of many years.

This system is used with state and local law enforcement to try to find matches to images from surveillance cameras and bank security footage. Most images are criminal mugshots. About 20 percent are civil images. The FBI started using these images without first having a privacy policy framework in place.

The other way is the face analysis team. It makes use of images in its own system, and other images from agencies, and Department of Motor Vehicles (DMB) databases in 16 states. They query many systems to try to identify images. The estimate is they have access to 411 million images to be used for their own investigative purposes. This access pre-dated any privacy framework for these images. The PIA was not updated until fall 2015. The face services side never had a privacy impact assessment (PIA). The SORN for these systems was not issued until this spring.

Ms. Levin: Was this system the same that was seeking exemptions from the privacy act?

Ms. Maurer: Yes, it is the same system. We felt pilot programs should not be exempt from privacy protections. They were searching on the photographs and images of Americans. There are real world consequences when the FBI shows up at your door.

The FBI was seeking exemptions on the same accuracy basis used by law enforcements. The GAO made recommendations to the FBI to do better. The GAO received recommendations from the Department of Justice to conduct a review as to why these programs existed so long without PIAs. DOJ felt what existed was sufficient based on the 2008 rules. The GAO disagreed and they still stand on their recommendations. The FBI did not do a good enough job. It is hoped the press will be the catalyst for change on this issue in the near term.

They have the policy in place not to do privacy impact assessments, but they spoke of their work reviewing PIAs. There was information from the SORN that it was ok. We feel the FBI did not do sufficient testing for accuracy before deploying to market. They did not do a good test of false positives. The DOJ did not stand by that recommendation. However, GAO does.

The vast majority of images the FBI has access to do not belong to the FBI. The FBI should take steps to make sure the information they are receiving is accurate. DOJ disagreed and did not feel the FBI had the authority or capability to do that. It is felt the FBI should accept all this information as accurate merely on faith.

Facial recognition technology must be accurate and the privacy rights of Americans respected. A number of states have opened DMV images to the FBI for scanning purposes. People should know this is happening.

The bi-annual Federal Information Security Management Act (FISMA) report will be issued next in 2017. This report focuses on 2013-2014 data. Most agencies continue to have significant weaknesses in the five areas focused on by the report. Most of the agencies had significant weaknesses.

Why are there weaknesses? What are the root causes? One of the main findings of the report was the risk management process. The central focal point must have access to management. Security should be viewed as an enabler to accomplish the mission not a barrier. Tony Scott, the Federal CISO, has pushed forward many of these issues. There are many things going on to try to improve these areas. DHS has been improving EINSTEIN, etc.

Public Participation

There were no public participants.

Internet of Things – Part II

Angelos Stavrous, Ph.D., Professor, Computer Science Department, Volgenau School of Engineering, George Mason University
Jeffrey Voas, Computer Scientist, Computer Security Division, ITL, NIST
Joseph L. Hall, Ph.D., Chief Technologist, Center for Democracy & Technology

Jeffrey Voas, Computer Scientist, CSD, ITL, NIST

This work started about two years ago. There was a lot of buzz about the security challenges of the internet of things. How do we think about the security challenges of an IoT system? In some ways, security is fundamentally no different than other IT systems, in other ways it's really different. They began with the question, what is the IoT? There is no actionable, and universally accepted definition of IoT. We're using the term, "this technology" instead of "IoT". "This technology" includes sensing, communications, computing, and some actuation. These four principles describe a NoT (network of things). For internets of things, there is no singular IoT, and it is meaningless to speak of comparing one IoT to another IoT.

An IoT is not measurable or boundable. A NoT is boundable, and all the other principles apply. It's bounded in the conceptual sense of people, process, and things. NISTIR8063 Primitives and Elements of Internet of Things is still in draft. It is awaiting final signatures. We anticipate it will be available in June or July.

There are five basic primitives:

- Sensor – Handles the sensing. There will be an incredible wealth of information.
- Aggregator – It is the big data engine inside the NoT. It merges all the data down to something that is digestible and usable. The aggregator makes the data actionable.
- Communication channel-governs how data is transmitted, it does the talking to other things.
- eUtility - External utility, software, hardware, or service, that feeds information into the workflow of a network of things. A human can act as an eUtility.
- Decision Trigger creates the final results needed to satisfy the purpose, specification, and requirements of a specific NoT. It has to do something of value for the user.

For each primitive, there are basic properties, assumptions, recommendations, and statements. There are 29 about sensors. There are 11 about aggregators. The actuation trigger acts based on data received in the NoT.

What has changed? Scalability, heterogeneity, lack of trust in the pedigree of data. Elements are other things to be considered: The environment, the operational profile for the NoT; Cost or expenses incurred by a NoT; Geographic location; owners, either a person or organization that owns the NoT; and a snapshot (time). The synchronization of things makes the network do what it's supposed to do. The snapshot can be vulnerable to tampering or malicious intent. These are big security challenges. This is a research work in progress.

To summarize, the internet of things is a catalog of technologies. Networks of things are far more measurable. Primitives and elements offer a "science". The goal is to someday measure trustworthiness. The state of a network

of things at a particular snapshot is a combination of primitives and elements for each asset in the first set, when applicable.

This model has been helpful in answering the basic question of what is the IoT. There has to be a scientific understanding of what the pieces and parts are. Is trustworthiness an example of a value, or a combination of all the other values? The term is being used loosely here, and is not to be taken as a formal definition. What is "pedigree"? It is the DNA of the information, the provenance of the information.

Angelos Stavrous, Ph.D., Professor, Computer Science Department, Volgenau School of Engineering, George Mason University

We have blockchain based protocols in IoT systems, but this is not a practical usage. The biggest problem in the internet of things is lack of encryption. IoT is not served well by the crypto community. They are focused on large groups. Individuals do not matter; large groups of data matter. In a home system automation setting, an attacker can introduce a soft AP or sensor with the same characteristics as the system. There is no authentication, no encryption or badly implemented, almost certainly a poorly designed system. We really need data integrity and authentication, security against single nodes, and handling sleep and power-off capabilities. Encryption is a plus, but not required.

Blockchain is at the heart of public data. In terms of health data, confidentiality may be less important than data integrity. There are short term and long term implications. The core of Blockchain is that parts of the data can be made available. Can we design blockchain protocols for IoT? An example is a home automation system. Attackers can introduce a new element to the system and collect all the data. Unsecured transmissions can be taken by attackers.

What do we really need? Dynamic but verifiable group membership, authentication and data integrity, secure against single node key leakage, lightweight operations A blockchain is a public distributed verifiable cryptographic ledger. All participants in the blockchain are able to "read", etc. We need to replace public key with hash functions. What we care about is the input. Data must be verifiable by all for a period of time. Ledger size is an issue. Group signatures can also be used. We need to replace PKI with hash based signature.

There are a lot of people struggling with the definition of IoT. The paper being discussed today is timely in this context. There is promise and peril in the IoT and the NoT. There is an amazing promise to the NoT. There are many amazing uses of sensors. There are also risks. User awareness and digital hygiene are constant problems. IoT devices put other devices on the network at risk. It is a good idea to use the guest network for IoT devices.

We are bringing together information from a number of groups, the IEEE and others. The American National Standards Institute (ANSI) and the International Business Machines Corporation (IBM) are working on a secure cloud for IoT. There are a number of grass roots efforts going on. The Federal Trade Commission (FTC) has done a significant amount of enforcement on devices in the market. There is a bill to set up a study group to monitor the internet of things.

Negative testing needs to have more emphasis. Credibility gap also exists with consumers.

Board – Wrap-up

There was a quorum present for Board business with Board members present and on the phone (Dave Cullinane and Gale Stone).

1. Ms. Sokol is leaving her position as ISPAB Designated Federal Officer (DFO). The Chair expressed appreciation for Ms. Sokol's hard work.
2. The Chair noted great topics were mentioned. Board letters for the following topics were suggested:
   a. Federal Privacy Council – write a letter covering privacy issues; PCLOB (recommending vacancies are fully staffed and the council is empowered to fill its role as authorized by Congress),

   b. FBI facial recognition, regarding the lack of consideration of accuracy in the technology and viewing pilots as exempt from privacy requirements;

   c. Transition planning on security priorities – The need to continue momentum on security, and recommendations to the Commission on Enhancing National Cybersecurity.

   d. Ms. Sokol sent out a Doodle request for input from the Board on future meeting dates. The Board will vote on the minutes and meeting dates during a call to be tentatively scheduled for the second week of July.

Ms. Levin moved to accept the motion to draft letters on the above topics. Mr. Greene seconded the motion. The motion was approved. Mr. Boyer will create the initial drafts for consideration by the Board. The drafts will be emailed and reviewed by the Board and approved when the language has been agreed upon. The Chair estimates six weeks for the process.

Openings within the Board – Patty Hatter from Intel will be joining next meeting. There is one other open spot at the present time. NIST is required to send a Federal Notice inviting the public to apply.

Future Meeting Topic: Invite Peiter Zatko (Mudge) to come to the next ISPAB meeting to discuss what his company is doing.

Board RFC Response: The Board may want respond to the request for comments on the internet of things.

Future Discussion Topic: Future Board discussion on prevention and focus on threats (Ron Ross).

Adjournment 12:28

The meeting was adjourned at 12:28 p.m.

ANNEX A

List of Participants

Last Name	First Name	Affiliation	Role
Annie	Sokol	NIST	DFO
Alam	Shireen	Symantec	Presenter
Aqua	Jocelyn	DOJ	Presenter
Badger	M. Lee	NIST	Presenter
Bailey	Leonard	DOJ	Presenter
Barrett	Matthew	NIST	Presenter
Beuse	Nathaniel	NHTSA	Presenter
Black	Paul	NIST	Presenter
Chen	Lily	NIST	Presenter
Felten	Ed	OSTP, White House	Presenter
Gilmore	Michael	GAO	Presenter
Groman	Marc	The White House	Presenter
Hall	Travis	NTIA	Presenter
Jagielski	Karen	FTC	Presenter
Kuhn	Rick	NIST	Presenter
Maurer	Diana	GAO	Presenter
Medine	David	PCLOB	Presenter
Moody	Dustin	NIST	Presenter
Morris	John	NTIA	Presenter
Plocher	David	GAO	Presenter
Polk	Tim	OSTP, White House	Presenter
Regenscheid	Andrew	NIST	Presenter
Ross	Ron	NIST	Presenter
Rudolph	Trevor	The White House	Presenter
Scholl	Matt	NIST	Presenter
Shabat	Matt	DHS	Presenter
Stine	Kevin	NIST	Presenter
Todt	Kiersten	NIST	Presenter
Wilkerson	Jessica	House Energy and Commerce	Presenter
Drake	Robin	Exetergov	Staff
Salisbury	Warren	Exetergov	Staff

Last Name	First Name	Affiliation	Role
Dodson	Donna	NIST	Visitor
Evans	Alison	Palo Alto Networks	Visitor
Kovacs	Kelsey	Symantec	Visitor
Marron	Jeff	NIST	Visitor
Norton	Paul	ID Experts	Visitor
Romine	Charles	NIST	Visitor
Suh	Paul	DHS	Visitor
Higgins	Joshua	Inside Cybersecurity	Visitor/media
Noble	Zach	Federal Computer Week	Visitor/media
Otto	Greg	Fedscoop	Visitor/media
Ravindranath	Mohana	Nextgov	Visitor/Media
Rockwell	Mark	Federal Computer Week	Visitor/media
Steinstein	Aliya	GovEx	Visitor/media