

INFORMATION SECURITY AND PRIVACY ADVISORY BOARD

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

MEETING MINUTES

March 3 and 4, 2021

Virtual Meeting Platform: BlueJeans

<p><u>Board Members</u> Steve Lipner, SAFECode, Chair, ISPAB Dr. Brett Baker, NRC Douglas Maughan, NSF Akilesh Duvvur, IBM Jessica Fitzgerald-McKay, NSA Brian Gattoni, DHS Marc Groman, Privacy Consulting Arabella Hallawell, NETSCOUT Systems Phil Venables, Google Cloud</p>	<p><u>Board Secretariat and NIST Staff</u> Matthew Scholl, NIST Jeff Brewer, NIST Caron Carlson, Exeter Government Services LLC Warren Salisbury, Exeter Government Services LLC</p>
---	---

Wednesday, March 3, 2021

Welcome and Opening Remarks

Steve Lipner, Chair, ISPAB, Executive Director, SAFECode

ISPAB Chair Steve Lipner, executive director of SAFECode, opened the meeting at 10 a.m. ET. Noting that there were many attendees in the audience, he provided some history of ISPAB:

- ISPAB was established by the Computer Security Act of 1987 to advise the U.S. government on issues of non-national security systems, including issues related to unclassified systems. The board, which first met in 1989, is made up of a dozen members plus the chairperson. Four members are from government agencies, four are from IT vendors, and four are from non-vendor organizations in the private sector.
- ISPAB is subject to the Federal Advisory Committee Act (FACA). The advice the board provides can take the form of letters to the heads of government departments and agencies.

The Chair invited board members to introduce themselves:

- Arabella Hallawell, Vice President, Strategy and Communications, NetScout Systems
- Akilesh Duvvur, Vice President, Worldwide Cloud Platform Product, & Offering Management, IBM
- Jessica Fitzgerald-McKay, Co-Lead, Center for Cyber Security Standards (CCSS) National Security Agency
- Philip Venables, CISO, Google Cloud
- Douglas Maughan, Head of Office of Convergence Accelerators, National Science Foundation

- Marc Groman, Principal, Groman Consulting, Adjunct Professor, Georgetown University Law Center:
- Brian Gattoni, Chief Technology Officer within the Cybersecurity, and Infrastructure Security Agency (CISA), Department of Homeland Security

Matt Scholl, Chief of the Computer Security Division in NIST's Information Technology Laboratory (ITL), reviewed the FACA, the board's role in providing consensus advice to the government, and how the board members are vetted. The role of the audience is to observe and listen but not engage with the board or speakers. There was a session scheduled at 3:30 p.m. ET for open, public dialog with the board.

Welcome and ITL Update

James St. Pierre, Acting Director, ITL, NIST

Mr. St. Pierre, ITL Acting Director, began his presentation with an overview of NIST leadership.

- NIST is in a transitional phase as the new Administration was settling into place. The Secretary of Commerce nominee, Governor Gina Raimondo of Rhode Island, was confirmed the previous day and was expected to be sworn into her new role later in the day.
- Jim Olthoff is currently performing the non-exclusive functions and duties of NIST Director. Other recently named leadership positions include: Chuck Romine, acting chief of staff; Eric Lin, acting associate director for lab programs; Stephanie Hooker, acting director of the Material Measurement Lab; and Joannie Chin, acting director of the Engineering Lab.

ITL Initiatives Update:

- Inclusive Language: As a result of a letter ISPAB sent to NIST in 2020, the agency updated the NIST Technical Series Publications Author Instructions. Additionally, the American National Standards Institute is working on gender responsive standards.
- Privacy Framework: Released on January 16, 2020, the Privacy Framework has garnered 31,000 downloads, and there are more than 40 implementation resources. More than one-quarter of respondents to a survey by the International Association of Privacy Professionals and FairWarning said they use the Privacy Framework. The majority of users are in the United States, Canada, the United Kingdom, India, and Brazil.
- Other noteworthy updates:
 - Draft FIPS 201-3, Personal Identity Verification (PIV) of Federal Employees and Contractors
 - Updated USG IPv6 Program in support of new federal IPv6 initiatives
 - Validating cryptographic modules to FIPS 140-3, security requirements for cryptographic modules
 - DevSecOps and Zero Trust Architecture virtual conference
 - Draft NISTIR 8312, Four Principles of Explainable Artificial Intelligence
 - Supplemental Materials for SP 800-53 Rev 5, Security and Privacy Controls for Information Systems and Organizations
 - ISO/IEC TS 27110:2021: Information Technology, Cybersecurity and Privacy Protection – Cybersecurity Framework Development Guidelines
 - There are four new NIST publications to help ensure that government and IoT manufacturers have a common understanding of IoT cybersecurity. [See details below.]
 - NIST guidance will help address challenges raised in the IoT Cybersecurity Improvement Act of 2020.
- SolarWinds Cybersecurity Breach: This breach demonstrated the need for strong trust roots in IT and cybersecurity, such as trusted software updates and digital signatures.

- Executive Order 13905 - Responsible use of Positioning, Navigation, and Timing (PNT) Services: NIST was charged with developing a cybersecurity profile to help organizations make deliberate, risk-informed decisions on their use of PNT services. Also, NIST will offer a time service over optical fiber lines as an alternate source of precision time.
- Encryption
 - Post Quantum Cryptography: Round Three selections have been completed.
 - Lightweight Encryption: The fourth Lightweight Cryptography Workshop held October 19-21, 2021.
- Trustworthy AI
 - Bias in AI Report will come out in this spring.
 - TREC Fair Ranking track will investigate bias in AI.
 - NIST is planning an international campaign to coordinate efforts in AI standard development with the Office of Science and Technology Policy (OSTP).
 - Explainable AI: The AI workshop on January 26-28, 2021, drew 900 participants from 24 countries.
 - U.S. Strategy for Resilient Manufacturing Ecosystems through AI: December 2-4, 2020, hosted by the National Science and Technology Council, Subcommittee on Advanced Manufacturing and Subcommittee on Machine Learning and AI
 - Face Recognition Accuracy: There has been improvement in recognizing masked faces. NIST recently analyzed 65 newly submitted algorithms using a database of 6.2 million images with masks digitally applied.
- 2022: Celebrating 50 years of cybersecurity research at NIST

Discussion:

- Mr. Groman noted that there a dozen or so privacy bills on Capitol Hill, and almost all of them touch on trustworthy AI, explainable AI, fair AI, etc. How does NIST interact with policymakers and legislators?

Mr. St. Pierre said NIST works with legislators through its legislative affairs office. NIST experts are often called on to provide technical advice.
- Ms. Fitzgerald-McKay asked how NIST draws the line between establishing standards to facilitate information sharing and leaving things to the implementers so as not to stifle innovation.

Mr. St. Pierre said that it is NIST's job to promote innovation. There is ongoing work on differential privacy, and they are looking at how an organization can know what information can be shared without creating privacy concerns.
- Mr. Groman said the work on AI or any other particular discipline falls under the larger risk management framework. The use of the output from AI will be different from agency to agency and application to application. Everything should have a risk assessment element.

Mr. St. Pierre said a large part of the focus is working toward a risk framework for AI. It is not a one-size-fits-all approach.

Elham Tabassi, ITL Chief of Staff, said the risks are different for different applications and uses. NIST is working on risk management beginning with the development of a taxonomy of risk. There needs to be a shared understanding of terms. For example, the computer science community might think of explainable AI differently from the way users thinks of it. Once there is a taxonomy and shared understanding of terminology, they will move on to metrics. They have to work on vertical components as well as.

- The Chair asked if the new standard for PIV cards will adapt to post-quantum encryption. Is there a transition plan in place for that move?

Mr. Scholl said yes, and they will share more information during the panel later in the afternoon.

- Ms. Hallawell asked whether there are any other big projects on horizon, such as 5G or industrial control systems.

Mr. St. Pierre said ITL is working closely with the Communications Technology Lab on a big project on 5G cybersecurity engagement. They are also continuing work on guidance for Industrial Control Systems (ICS) cybersecurity.

Kevin Stine, Chief of ITL's Applied Cybersecurity Division, added that NIST has a long-standing role in developing guidance for ICS cybersecurity. They are planning to release an update to NIST SP 800-82, *Guide to Industrial Control Systems (ICS) Security*, likely by the end of the fiscal year. This is of interest to a number of sectors, including those that have experienced recent cybersecurity challenges. There will continue to be an emphasis on software security and supply chain risk management activities, and they are beginning to formalize the approach and resources in those areas. They can provide more details later in the meeting and at future meetings.

The Chair recessed the meeting for a 15-minute break.

SolarWinds Forensic Brief

Jay E. Gazlay, Technical Strategist, Cyber & Infrastructure Security Agency, Department of Homeland Security

Mr. Scholl introduced Jay E. Gazlay, Technical Strategist for the Cyber & Infrastructure Security Agency (CISA) at the Department of Homeland Security.

Mr. Gazlay provided a brief overview of his background and a review of the recent SolarWinds supply chain attack.

- Key Takeaways from the Attack:
 - The adversary that hacked SolarWinds was patient and resourceful. We need to remember that adversaries are thinking in the long term, year over year.
 - The adversary is exploiting weaknesses in our supply chain and identity management systems.
 - Follow-on actions on objectives are very difficult for many organizations to identify. Many agencies did not detect the hacking activity for months. It is really hard to appropriately understand what is going on with identity, and the targeting of incident responders adds new complexity.
- Timeline Overview:
 - 9/4/19: Threat Actor (TA) accesses SolarWinds.
 - 9/12/19: TA injects test code into SolarWinds build process and begins trial run to see if code is detected.
 - 2/20/20: SUNBURST malware is compiled and deployed.
 - 3/26/20: Hotfix 5 DLL made available to customers.
 - 6/4/20: TA removes malware from build VMs.
 - 12/12/20: SolarWinds notified of SUNBURST after it was discovered at Mandiant by a helpdesk technician who noticed the addition of a new Multifactor Authentication token.
 - 12/14/20: SolarWinds files 8-K with the SEC and notifies shareholders and customers.

- 12/15/20: SolarWinds releases software fix.
- 12/17/20: US-CERT alert is issued.
- 1/11/21 New findings related to SUNSPOT are released
- CISA's takeaway:
 - Identity is everything now: Adversaries gained and used legitimate credentials in the attack and operated for months in many environments without detection.
 - TrustStore and IDM (Identity Management) compromises are excellent targets that adversaries are exploiting. This becomes more pernicious as we move further into the cloud infrastructure where all that matters for access is the assertion that you are who you say you are. We need to take care in how identity structures are managed and monitored. With dependence on the Cloud environment, new guidance on identity protection is needed.
 - Behavioral analysis techniques are required to tell that identity compromise is occurring. For example, a user in Maryland logging in locally to a machine in California.
- Detection Opportunities:
 - Detecting a supply chain compromise of this nature is beyond most organizations' capabilities.
 - Network baselining and abnormal behavior analytics are instructive.
 - User behavior abuse
 - Impossible logins (different locations at the same time)
 - Secure Assertion Markup Language (SAML) abuse: See a lot of SAML assertions that are only good for a minute or for 2 years.
 - See US-CERT Advisory AA21-008A – Detecting Post Compromise Activity in Microsoft Cloud (Sparrow – <https://github.com/cisagov/Sparrow>)
- Key Questions to Ask:
 - Do you know who can be trusted? When did you last validate?
 - Do you have visibility into your hosted/cloud environments? Can you see all authentication attempts?
 - If your main network was compromised, can you continue to operate? Do you have backup communications?
 - When did you last exercise your Disaster Recovery plan?

Discussion:

- Mr Venables asked if there is anything NIST could be doing related to standards to shift detective and protective behavior, given what we've learned with SolarWinds?

Mr. Gazlay said we need to have basic configurations across all of our devices. There is a small number of extremely high-value configurations that can be broadly deployed. One example is RFC compliance on firewalls so you know traffic is being formed in a normalized fashion. We should work together with the vendor community to prioritize detection and security of identity on a small number of configurations. That is the type of partnership we should have with industry. The National Checklist Program and the use of SCAP to automate secure configuration management can be important.
- Mr. Duvvur asked about fourth-party risk and scenarios where you are leveraging a Software as a Service (SaaS) provider, for example. You have the underlying infrastructure on a cloud, so there is a chain of complexity that is built in. Are there guidelines for managing or ameliorating that situation?

Mr. Gazlay said it is an emerging topic that is generating a robust and active dialogue. There is a lot of value in things like SaaS that enforce an upgrade cycle on large-scale entities but threat actors are going after identity. It's a different risk calculus with significantly sophisticated threat actors.

- Ms. Hallawell asked about secure configurations and working with vendors, particularly in near real time. For the majority of companies, threat hunting is still not scalable. How do we train the teams to be more successful? How could the industry or agencies help with automated tools?

Mr. Gazlay said he would love to see standard inputs and outputs for automating IOC detection and different ways of doing that. It would be great if industry and NIST could partner on standards around that and how the apps work so they could work across multiple different operating systems and platforms to perform the first tier of problem identification. In terms of educating Security Operations Centers (SOCs) on how to do that, he has frequently made the point that the best thing to do would be to stand up a Slack instance and drop everybody on the SOC onto it. This issue speaks to human communication. Being able to communicate in a way that bridges those really high-skill people and new people would be useful. Often SOC contracts have been labor dollar heavy and not sophistication heavy. How do we shift the narrative to function specific and innovation oriented?

- Mr. Groman commented that in the federal government, decisions to move forward with a new technology do not always include good risk assessments. For example, there could be an executive order that says every agency needs to adopt a new technology. If we are duct taping legacy systems to something new, we should not be shocked when an incident occurs. The Risk Management Framework, OMB Circular A-130, all of that in theory requires that the head of the agency is accepting the risk that, but who knows what they're signing off on or who's making the decision.

Mr. Gazlay said the point is well taken. The discussion around where risk is being accepted and who is accepting it is valid. The concern is seen in the private sector as well. People have accepted a lot of risk to identity they did not know they were accepting. The language from cybersecurity professionals is often so stilted and unintelligible to business executives that they cannot be blamed for making the wrong risk assessment.

Mr. Groman agreed that it is a concern that the wrong people are accepting risk or they don't appreciate it or understand it. When a breach occurs and data is compromised, the agency where the hack was initiated does not necessarily feel the consequences, and we're not addressing that.

- The Chair asked whether agencies have suitable guidance that they are not following or whether the guidance is not helpful. Did agencies not follow the Risk Management Framework, or did they follow it and it didn't help? What should the government be doing?

Mr. Gazlay said that with regard to identity, the guidance should be updated to include the Cloud. It is extremely difficult to follow all of the guidance because the infrastructures are indefensibly big. There is no standardized guidance or program to recover from a nation state level adversary. NIST 800-160 Volume 2 has been the north star as they talk to agencies as they modernize – the idea of building resilient systems from the start that can easily take that first punch. A lot of agencies can't survive first contact. Making guidance actionable is hard.

- Ms. Fitzgerald-McKay asked about the role of automation in defense. Is there a role to be played in preventing the next SolarWinds?

Mr. Gazlay said that we can only solve machine-to-machine problems with machine-to-machine solutions. The entirety of our daily business is touchable via APIs, and we don't understand the risk. Only through automation on those fronts can we get the problem solved. He believes in the SCAP-type solutions. He believes in security containers and asking, as you build infrastructures in the Cloud, how do you automatically build these detection capabilities into your infrastructures?

- The Chair noted that the National Checklist Program had not been updated in a long time.

Mr. Gazlay said there has been a little activity, but programs like that need to be more prevalent and more used. Vendors know how to protect their products best. If he wants to know how to secure a product, he calls the product security team. In the future, we would want to get to the point where vendors are releasing some configuration guidance in a normalized format that is machine readable and answerable so that people can make educated risk decisions at scale across distributed and federated infrastructures.

The Chair said he wanted to return to this topic in the afternoon and discuss whether there is advice that the Board can provide.

The Chair recessed the meeting for a 1-hour lunch break.

IEEE Paper: It Lurks Within: A Look at the Unexpected Security Implications of Compliance Programs

Dr. Michelle Mazurek, Computer Science Department, University of Maryland

Rock Stevens, University of Maryland

Dr. Josiah Dykstra, Technical Fellow in the Cybersecurity Collaboration Center at the National Security Agency

Mr. Scholl introduced Dr. Michelle Mazurek and Dr. Josiah Dykstra.

Dr. Mazurek began her presentation with an overview of compliance standards. She noted that they are often used as checklists and can have limited effect even when used correctly. An example is seen in the recent \$10 million fine that North American Electric Reliability Corp. (NERC) issued for security lapses. In her study, Dr. Mazurek's team looked at what might still go wrong even when an organization follows standards correctly.

- Main Components of the Study:
 - Auditing procedure: Multiple researchers examined standards and identified security issues, categorizing by root cause and risk level.
 - Risk assessment procedure: Researchers looked at the severity and probability of the security issues.
 - Expert evaluation procedure: Researchers sent a subset of the issues found to domain experts and asked if they 1) had seen them; 2) had not seen them but consider them plausible; or 3) do not consider them plausible. They asked the experts to challenge their assumptions and provide any additional context for how the standard is used.
 - Disclosure process: They informed authorities, standards bodies, and users.
- Standards Examined:
 - IRS P1075 – governs handling of tax payer data: Evaluated by compliance officers from NYC
 - PCI DSS – governs handling of credit card data: Evaluated by compliance auditor

-
- NERC 007-6 – governs security of electric power grids: Evaluated by a professional who had contributed to development of the standard
 - FedRAMP – examined separately following the study
 - General Findings:
 - IRS P1075 has the most issues. NERC 007-6 has the fewest.
 - Experts generally agreed with the research findings. In some cases they rejected a concern by arguing that, in real life, people don't rely on just one standard.
 - Four Vulnerability Categories:
 - Ambiguous specification: Language can be interpreted differently by different readers. This often includes optional controls, passive voice, and a lack of deadlines. Examples:
 - IRS P1075: Access control policies must be evaluated every 3 years, but by whom? What does the evaluation consist of?
 - PCI DSS: All issues identified during a pen test must be addressed, but how and when?
 - Data vulnerabilities: Insufficient data protection. Examples:
 - IRS P1075: Covers only data originating from IRS, not duplicates obtained from other sources.
 - PCI DSS: The term “sensitive data” does not include passwords, SSNs, DoBs.
 - Unenforceable requirements: Requirements cannot be enforced in practice and are often contradictory. Examples:
 - IRS P1075: Requires multiple forms of physical security to protect access, but also authorizes telework and remote access.
 - Under-defined processes: Missing steps that create gaps in security. Examples:
 - IRS P1075: Mandates network component inventory but never establishes a baseline.
 - Disclosure Attempts:
 - Common Vulnerabilities and Exposure (CVE) operators said they only work on software issues that can be patched.
 - National Vulnerability Database (NVD)/MITRE said this sort of vulnerability was out of scope
 - Department of Homeland Security (DHS) at first expressed interest and then issued an ultimatum to cease communications.
 - PCI Working Group accepted some changes.
 - NERC said it would review the findings.
 - NIST/National Cybersecurity Center of Excellence (NCCoE) said there will be updates in ongoing revisions.
 - FedRAMP Study:
 - Background: The pandemic led to massive adoption of videoconferencing, including Zoom. CitizenLab identified several security problems in Zoom, and many government agencies banned it in response. However, Zoom for Government was already FedRAMP compliant.
 - The researchers evaluated FedRAMP, based on NIST 800-53 Rev 4. They used the same audit method and auditors as in the full study, but no external evaluation. They found new but overlapping vulnerability categories:
 - Ambiguous specification: They found 11 instances without a time specification (e.g., privileged execution must be audited, but how often?); 10 authentication issues (e.g., weak passwords, poor crypto); define-your-own requirements (e.g., access control procedures; what “sensitive data” means).
 - Data Vulnerabilities: They found no protection requirements for crypto keys; insider threats.

- Obsolete references: They found references to FIPS I40-2, which was replaced in September 2019 and a password expiration policy that has since been superseded; FedRAMP itself was updated in August 2018.
 - On the plus side, FedRAMP did better than the other three standards. It involved collaborators, lessons learned, and requests for comment. There was improved training and evaluation for third-party auditors.
- Underexamined Threat Models:
 - Nation-state access: Where is critical information stored? Can a government mandate access?
 - Aggregation and monetization: Modern software is designed for tracking and aggregation - how does this fit in government products?
 - Cyberphysical systems need to be better accounted for.
 - Security of the security tools: FedRAMP doesn't consider a threat model of the tools themselves. There are no checks and balances.
- Recommendations:
 - Check-listing is inevitable, so plan for it.
 - Minimize ambiguity: Some amount of ambiguity is inevitable, but it might be possible to minimize ambiguity or put in additional procedures that set bounds for ambiguity.
 - Design for supplementation: A lot of domain experts talked about their awareness of potential issues and explained that that's why the compliance standards are just a baseline - but now you have created controls that are not part of a compliance regime and are not audited in the same way.
 - Continuous updating: Standards very quickly can become outdated. Try to build flexibility into the process - maybe open requests for comment or open disclosure processes. Maybe more frequent updates, but with a grace period built in.
 - Consider more threat models, such as insider threats and security tools.

Discussion:

- Mr. Groman asked about the assumption that anyone using the standard knows it's not a checklist but something that evolves with the organization and the technology.

Dr. Mazurek said in some ways that is at the core of the problem. The standards need to be treated as a living thing, but that doesn't happen in practice. They just get used as checklists in a non-thinking way. We need to try to figure out bridges to make it work better.

Dr. Dykstra said some standards have a degree of reasonableness. There can be a spectrum of compliance, but it probably comes at a cost.

Mr. Duvvur said check listing is seen everywhere. There is confusion about different terms. Is there a way for better end-to-end education?

- The Chair said often the user is focused on getting a business capability deployed or satisfying his or her boss with the latest feature. The compliance program is done in a way that interferes as little as possible with the business at hand. Getting the business manager to understand that security is something they have to do is a cultural change issue. It takes a lot of personal impact to make that happen.
- Ms. Fitzgerald-McKay asked about adding configurations and addressing the low-hanging fruit to move the bar forward.

Dr. Mazurek said these questions go to useable computing issues.

Dr. Dykstra said that getting the community to take the right security steps involves human problems.

- Mr. Groman said that we have not properly incentivized security in our country - in the public or private sector. There has to be a mandated incentive to prioritize accountability.

The Chair said he has seen a couple of instances of private sector organizations that really got a shock that resulted in a cultural impact. It takes a significant shock at a high level and a message from a cultural leader to effect change. It's a rare event.

The Chair thanked the speakers and recessed the meeting for a 9-minute break.

Open Source Software Security Study

Dr. Frank Nagle, Harvard Business School

Dr. David A. Wheeler, Director of Open Source Supply Chain Security, Linux Foundation

The Chair introduced Dr. Frank Nagle of the Harvard Business School and Dr. David A. Wheeler, Director of Open Source Supply Chain Security at the Linux Foundation.

- Overview of open source software (OSS):
 - Licensed to users with the freedom to run the program for any purposes, study and modify the program, and freely redistribute copies of original or modified program
 - Under U.S. law and regulation, it is commercial software
 - Typically developed and released through a trusted repository with the user as developer
 - OSS licenses enable worldwide collaboration
 - Well-run OSS projects seek to nurture collaboration
 - OSS is not “no cost,” but cost-sharing and collaborative review often make it low cost
 - The Open Source Security Foundation (OpenSSF), established in August, 2020, focuses on improving the security of OSS. Members include GitHub, Google, IBM, Intel, Microsoft, Red Hat, Uber, and VMware.
- 2020 FOSS [Free and Open Source Software] Contributor Survey Overview:
 - Dr. Nagle reviewed the study. They looked at respondents' demographics, current FOSS contributions, whether they are paid by their employers for their work on FOSS, motivations, time allocation, and possible incentives. For example, what levers can be pulled to nudge contributors in the right direction?
 - There are three main types of open source developers: paid maintainer, paid occasional contributor, and unpaid maintainer
- Demographics:
 - Approximately 1,196 survey respondents (¼ from U.S.; ¼ from Germany/France/U.K.)
 - 93.2 % Male
 - 3.17% Female
 - 35-44 years old - most prevalent age range
 - 51.65% receive payment for their FOSS contribution from either their employer or a third party. (63.8% of respondents in U.S. were paid; 15.8% of respondents in India were paid.)
 - Most respondents are employed full time
 - 37% work in the software development sector; 17% work in IT services
 - 39% received formal training in secure software development
- Employer's IP Policy related to FOSS contributions during free time:
 - There is still a high percentage of companies with an unclear policy or employees who don't know the policy.
- Motivations:

-
- 1st - They use open source and need the specific feature they add
 - 2nd - They enjoy learning
 - 3rd - It fulfills a need for creative, challenging, or enjoyable work
 - Last – They are paid to develop open source. Other low motivators were that it would advance their career or give them important peer recognition.
 - Time Allocation:
 - Security is low on both how contributors want to spend their time and how they actually spend their time.
 - Most beneficial contributions from an external group would be:
 - 1st – Help with bug and security fixes
 - 2nd - New features
 - 3rd - Financial help
 - Summary and Suggested Actions:
 - 74.9% of respondents are employed full time → Leverage the top three motivations for contributing (desire to learn, need for features/fixes, need for creative work)
 - o Recognize value of skills gained from FOSS contributions
 - o Support learning process for new contributors
 - o Balance creative and mundane tasks for all contributors
 - o Consider other financial support (e.g., security audits, travel, etc.)
 - 2.8% reported that contribution time is spent on security issues
 - o Fund security audits for critical FOSS projects
 - o Prioritize secure software development (SSD) best practices
 - o Make SSD training required for paid FOSS developers
 - o Incorporate security tools and automated tests into Continuous Integration (CI) pipeline
 - 48.7% are paid by their employer to contribute to FOSS
 - o Allay concerns over corporate involvement in FOSS through greater transparency and clear commitments
 - o Incentive paid contributors to dedicate time to mentoring new volunteer contributors
 - o Transfer FOSS projects to foundations w/ neutral governance
 - 17.5% said their employer had unclear FOSS policies
 - o Clarify policies on when and how employees can contribute to FOSS
 - o Promote contributions to FOSS projects' security improvements - through individual employee's or collaborative efforts (e.g., OpenSSF)
 - What can be done to improve OSS Security? Improve security practices while limiting the burden on contributors:
 - Identify vulnerabilities in existing code and propose fixes
 - Help modify CI pipelines to add problem-detecting tools
 - Audit critical OSS projects and develop patches
 - Rewrite portions/components prone to vulnerabilities
 - Contribute hardening measures
 - Require secure dev training for paid OSS developers
 - Use badging programs for secure dev practices
 - Ask influential OSS contributors to stress security
 - Partner with mentoring programs to incorporate security best practices
 - Simplify incorporating tools into CI pipeline
 - It is possible to improve OSS security:
 - OpenSSL post-Heartbleed: Linux Foundation Core Infrastructure Initiative (LF CII) founded and invested in OpenSSL, dramatically improved

- LF/CII audits – funded many security audits to proactively find/fix problems
- OSS-Fuzz/Fuzz onboarding – contractors fuzz OSS, report vulnerabilities to OSS projects
- Google/LF underwriting some Linux kernel security development efforts
- Ideas for Government:
 - Identify OSS critical to U.S. government and critical infrastructure and invest.
 - Require developers of U.S. government custom software to know how to develop secure software. There is a free course released by OpenSSF.
 - Encourage maturing and use of Software Bill of Materials.
 - Update/Eliminate the Vulnerabilities Equities Process.
 - Ask NIST to formally define reproducible builds, and move to requiring reproducible builds for high-criticality software.
 - Require U.S. ISPs to secure infrastructure. Require Resource Public Key Infrastructure (RPKI) to protect Border Gateway Patrol (BGP).
 - Fund NVD/CVE to proactively track vulnerabilities. Right now they depend on reporting, and that is missing a lot of vulnerabilities.

Discussion:

- Mr. Maughan asked if there was anything in the survey about the incentives and being funding to fix vulnerabilities. Was there a question about when someone knows they have vulnerabilities, how often to fix them, or if an external funding source would make a difference in the vulnerabilities that get fixed?

Dr. Nagle said the survey approached vulnerabilities in two ways, current and future. Financial incentives to conduct security audits and paying people to fix them is one thing; educating people to write secure code from the beginning is trickier. Perhaps people could be paid to take the courses. Also, employers could have a positive impact by tying secure code to bonuses.

Dr. Wheeler said the big problem is when there are vulnerabilities the developer is unaware of. Improving information that gets back to them so they can fix them is important.

- Mr. Maughan asked if there were questions in the survey about the utility and value of bug bounty programs.

Dr. Nagle said bug bounties were mentioned in one or two places in the survey, but they didn't ask about their value or utility.

Dr. Wheeler said bug bounties can be effective if they are implemented after you have done an audit, run tools, and hardened the system. However, if you're paying people to find defects that can be found in 10 minutes, it's a waste of money.

- Mr. Venables asked about getting people who use open source to install the latest updates. Are we looking at the end-to-end packaging, update process, etc.?

Dr. Nagle said they conducted a different study on that issue. They released a report on open source usage and are in the process of gathering data for an update. They will release details in the next month or two.

Dr. Wheeler added that automation is the only way to resolve this.

- The Chair said that getting developers trained on software security and secure development is an important measure. He was surprised to see that 39% of the respondents claimed to have been exposed to secure development training.

Dr. Nagle said the bar was low for defining security training.

The Chair recessed the meeting for an 11-minute break.

NIST Cybersecurity and Privacy Update

Matthew Scholl, Information Technology Laboratory, NIST

Kevin Stine, Information Technology Laboratory, NIST

Mr. Stine opened the presentation with a personnel update. Jeff Greene, NCCoE director, will be going on detail to the National Security Council (NSC). Natalia Martin was named NCCoE acting director.

Mr. Stine and Mr. Scholl provided a programmatic update, starting with an overview of NIST's cybersecurity and privacy priority areas. NIST is trying to approach priority areas in a strategic way, tackling current challenges while anticipating future challenges. There are nine strategic priority areas, and there are factors that cut across them, including human factors and usability, standards, the international community, and automation.

- **Enhancing Risk Management**

NIST views cybersecurity and privacy in the larger context of business activity. They want to improve the resources they produce to help integrate and strengthen ties between privacy and security within the broader enterprise risk management umbrella. They aim to help risk decision makers at an enterprise level understand and communicate risks. The enterprise could be at a department level or the U.S. government. Recent resources include:

 - NISTIR 8286: Focuses on integrating cybersecurity and Enterprise Risk Management (ERM), taking advantage of some common ERM tools, such as risk registers.
 - Draft NISTIR 8286A: Focuses on identifying and estimating cybersecurity risk for ERM – identifying processes and capabilities intended to be actionable.
 - Automation: They are trying increasingly to bring automation into risk management.
 - Just as an aside, on March 3, 1901, the National Bureau of Standards was established.
 - SP 800-53 Rev 5: NIST redoubled efforts to ensure that control catalogs are available to tools, Governance, Risk, and Compliance (GRC) screens, compliance checkers, and those who wish to use them in other formats. The data is available in multiple formats for automation. They anticipate that it will be deployed before the next ISPAB meeting.
 - Updates to SP 800-53A (assessment procedures) and SP 800-53B (baselines)
- **Privacy**
 - Privacy Framework: NIST recently celebrated the first anniversary of the Privacy Framework, and they are very happy with the uptake. A lot of new resources are cropping up, some developed by NIST and some developed by other organizations.
 - There is increased international interest in the Privacy Framework. They recently released Spanish and Portuguese translations.
 - Additional crosswalks have come out between the Privacy Framework and other resources, including the California Consumer Privacy Act (CCPA) and SP 800-53B.
 - They are looking at increasingly bringing privacy into NCCoE projects and would appreciate feedback from the board on focused privacy-related projects they could take on.
- **Strengthening Cryptographic Standards and Validation**
 - Round Three selections for post-quantum cryptography are underway. There has been nice work done in analyzing some of the finalists. Some good papers have come out. They will take a hard look at post-quantum mechanics, including factoring lattice implementations. It is nice to be dealing with a smaller set of algorithms because the community can be much more focused.

-
- They are also looking at the issues around crypto agility – key sizes, computational overheads of potential finalists and what that means for different implementations in things such as a Transport Layer Security (TLS) exchange or a decryption capability. In the past, NIST would just flip the switch through the transitioning of the document. They know from experience that crypto transitions are extraordinarily difficult.
 - Crypto and Identity Management
 - PIV Cards and FIPS 201: They updated the form factor of the PIV card in FIPS 201, which specifies the form and the data on it as well as the hardware security module. The current crypto specified for PIV cards is not quantum safe, so they reference a specific implementation of RSA or specific implementation of an Elliptic Curve Digital Signature Algorithm (ECDSA) for interoperability issues that might occur if they allow for any digital signature implementation on the cards. When they're done with the FIPS 201 update, they will start updating the downstream guidance that specifies the other card characteristics.
 - SP 800-78: Next year, they will start updating SP 800-78, which will specify crypto on the card. It will be a timing issue. They will probably complete the SP 800-78 updates before the post-quantum cryptography standards are finalized. There will have to be forward pointers in SP 800-78 identifying the crypto to give them the flexibility to enroll them once they're ready. They will be looking at the card space and computing space on the chips called for on the cards. They will also be looking at performance and interoperability.
 - Identity bound to a cryptographic token is extraordinarily important for trust. They will be looking at federated identity with the PIV card for the future. An Office of Management and Budget (OMB) memo outlines the future of identity and identity uses in the federal government. They are looking to provide guidance on a federated structure for the use of potential commercial tokens for exchanging information with the federal government. That will be a new set of guidance developed either in 2021 or 2022.
 - Mr. Maughan asked if that meant he could take his National Science Foundation (NSF) PIV card and NIST would recognize it and allow him on the campus. If you have a FIDO [Fast ID Online] token and you want to get information as a citizen from the Social Security Administration, how does the commercially provided token map to security environments? Mr. Scholl said that it could potentially be something that could allow a citizen to get their information from the Social Security Administration, for example. How does a commercially provided token map to risk environments?
 - SP 800-63 Rev 4 is in the midst of an update, and they plan to have a draft out for comment in the late spring.
 - There is an NCCoE project focusing on practical and actionable ways to prepare for crypto migrations. There is also a draft project description out for public comment addressing visibility challenges as organizations move toward TLS 1.3. The project description is intended to draw comment from the community in helping NIST define the scope. Once they have the scope, they will reach out to solicit industry partners and collaborators.
 - Transition from FIPS 140-2 to FIPS 140-3: Both programs are valid during the transition. In April there will be a workshop on good cryptographic entropy sources, looking at what is a good source of entropy and how do you test it?
 - They will continue working with industry and at the NCCoE on further automation of cryptographic assessments.
 - Awareness, Training, Education and Workforce Development
 - Just about every conversation touches on workforce, education, training. Workforce development makes up another priority area for NIST. Awareness and training cut across the

- entire portfolio. Within the last few months, a new National Initiative for Cybersecurity Education (NICE) Strategic Plan was released with a refreshed set of strategic goals.
- They continue to review and update the NICE Framework, focusing on developing competencies and associated tasks, knowledge, and skill statements. They expect to release a NISTIR on NICE Framework competencies for public comment soon.
 - Workforce is a dimension of every discussion and increasingly with sector-specific engagement. Increasingly, they are asked to consult on how resources can be adapted and tailored to plan for workforce needs.
 - Trustworthy Platforms
 - There is an increasing focus on software security, broadly speaking, to better understand core capabilities and the technical underpinnings needed to secure software, which includes more security-conscious development methodologies. They have hosted a couple workshops in the last month or two to inform their work in software security and development methodologies. They are also looking at technical capabilities around code signing, improving security, and development tools.
 - There is still research to be done in service mesh architectures' use of containers and policy enforcers in multi-cloud environments. There used to be a U.S. Government Configuration Baseline, and the lesson learned is that we were always two versions behind by the time they created it.
 - Other strategic priority areas include:
 - Trustworthy Networks, including projects on 5G, IPv6, and Zero Trust Architecture
 - Cybersecurity Measurement
 - Securing Emerging Technologies, including IoT

Discussion:

- Mr. Venables said it is correct to focus on automation, but you also have to re-imagine how things are done and not just automate existing things. Controls that are manually implemented today maybe should be included as inherent features in the platforms and not have to be subsequently configured. A slightly different way of looking at it is ambient control.

Mr. Stine agreed, adding that NIST uses the term “re-imagine” quite a bit.

- Ms. Hallawell asked about DevSecOps and the workshop NIST recently hosted. What about the latter stages when teams are monitoring and figuring out if there are anomalies? Is some of the guidance about the later stages of management of the solutions and environments?

Mr. Scholl said the workshop is online. There was a lot of discussion around where capabilities are in lifecycles. There is some guidance around very technical and tactical pieces, and they are continuing the work on it.

Public Comment, Summary of Day 1, and Board Discussions

Jeff Brewer, Designated Federal Officer for ISPAB, said he received no official requests from the public in writing prior to the deadline, but he did receive an email.

Mr. Scholl asked if the person who emailed was in attendance, and if so to raise their hand. Nobody came forward.

The Chair asked board members about topic areas that might warrant a recommendation:

- Mr. Maughan asked if the board should make a statement related to the OSS community and secure development.

The Chair asked if it would make sense to encourage OMB or the Department of Commerce to consider contractual clauses that encourage security training for people contributing to open source.

Mr. Maughan asked if NIST had ever made a statement about software development from the workforce perspective.

Mr. Stine said they had in the context of understanding the workforce needed to securely develop, for example.

The Chair said it is a broad workforce consideration, and the Board could go on record stating that it is a good and necessary idea. He said he would write a draft overnight.

Mr. Venable asked about an additional statement that encourages agencies that use open source to sign up for the OSS Foundation.

The Chair said the recommendation would probably encourage agencies using open source to join the OSSF or consider it a resource for the security of the software they rely on. Integrating broad developer security training is something NICE might pursue. He will write a paragraph on each recommendation to discuss on Day 2 of the meeting.

- Ms. Fitzgerald-McKay asked if there was a way for the board to help encourage secure configurations from software vendors.

The Chair asked Mr. Stine about secure configurations from vendors as part of the trustworthy platform initiative and whether it would be reasonable to recommend an effort to work with vendors to ensure secure configuration of out-of-the-box systems and platforms the federal government uses.

Mr. Stine said it would be useful to re-emphasize that.

The Chair asked Ms. Fitzgerald-McKay to draft a statement.

Ms. Fitzgerald-McKay asked what would be useful to capture.

Mr. Scholl said that what would be useful would be something that emphasizes the importance of the matter and requests that it potentially be made a priority for future capabilities.

- Ms. Hallawell asked whether they should consider something more specific to supply chain security. Is there a broader supply chain comment to make? Configuration is important in itself, but should it be part of a larger statement? Secondly, should there be a statement on automation?

Mr. Gattoni asked how automation does not become subject to the same fault modes it aims to remedy. If it's a fix-all, it becomes the next target. How can they insert that into the design process upfront?

The Chair said emphasis on automation is certainly important. He asked Mr. Gattoni to draft an initial statement and send it to Ms. Hallawell.

Mr. Gattoni agreed. His initial thought is that automation is code, so development principles still apply: Do it right.

Ms. Hallawell asked if it could be put on the agenda for a subsequent meeting. Automation and AI are often molded together in dangerous ways.

The Chair thanked the board members and speakers and adjourned the meeting for the day.

Thursday March 4, 2021

AI and NDAA requirements; NSC AI Commission Report

Elham Tabassi, Chief of Staff, Information Technology Laboratory (ITL), NIST

The Chair opened the meeting at 10:00 a.m. ET and introduced Elham Tabassi, ITL Chief of Staff.

Mr. Scholl announced that before the meeting adjourns, there will be a presentation by the Commerce Department's Office of General Council, which will provide the mandatory annual FACA ethics talk.

Ms. Tabassi began her presentation with an update on NIST's AI activities. AI is rapidly growing and transforming the world. The optimistic view is that it will change our lives for the better, but the pessimistic view is that it can be a threat to humanity. Of course, the truth lies somewhere in between. It can also have the potential to lead to some significant unintended consequences for individuals, organizations, or society.

Major advances in AI continue to drive a need for a universal understanding of its risks. It offers many possibilities, but has the potential to exacerbate the trend of rising inequality and threats. There hasn't been a concerted effort to describe a taxonomy of AI risks and how they can be managed. As we decrease risk, we increase trust. NIST has been engaging the public and private sectors about building blocks for trustworthy AI and associated measurement and standards.

- Trustworthy AI's Foundation
 - Identify building blocks or technical requirements: The technical requirements for core building blocks so far include accuracy, reliability, robustness, safety, security, privacy, explainability, and freedom from bias. Other factors, such as fairness, also need to be considered.
 - Establish concepts, terminology, and taxonomy: For each technical requirement, we need a clear and common set of vocabulary.
 - Metrics, evaluations, and benchmarks: Once we know what to measure and we have stable, agreed-upon definitions, we can work towards developing metrics and testbeds for benchmark and performance evaluations. A solid scientific foundation about what to measure and how to measure it is essential for developing clear, voluntary, globally relevant standards in an open, transparent, and inclusive manner.
 - Risk management: Allow developers and users to decide risk levels and manage interaction. There is no one-size-fits-all approach. Different uses within a technology have different levels of risk.
 - Governance: People are asking for a risk management framework for AI now.
 - Policy Considerations
- Core Building blocks of Trustworthy AI: We have all heard that AI systems are brittle, so the question is how to make sure the system is fail safe. The issue of bias is something we have to worry about. Depending on the definition, bias is not always a bad thing (statistical bias, e.g.). For example, young adults should pay more than older drivers for car insurance.
- So far, the bulk of the work NIST is doing is the foundational research focusing on how to measure enhanced, trustworthy, and responsible AI. NIST has produced a series of reports on the

different characteristics of trustworthy AI. Additionally, NIST is conducting user-inspired research applying AI to many different measurement problems.

- Four Principles of Explainable AI: Terms are being used differently by different communities. In August, NIST released for public comment a paper on the principles of explainable AI and received a very robust set of comments. They are working on finalizing the document.
 - Systems deliver accompanying evidence or reasons for all outputs.
 - Systems provide explanation that are understandable to individual users.
 - Explanation correctly reflects the system's process for generating the output.
 - The system only operates under conditions for which it was designed or when it has sufficient confidence in its output.
- Bias in AI Workshop (August 2020): They are finalizing an extended report, which will be released for comment soon. They plan to have other activities focusing on how to measure bias.
 - Need consistent terminology
 - Need standardized measurement of bias
 - Risk management of bias in AI
 - Understand datasets and algorithms within the context of use cases
 - Include a diverse range of scientific and other scholarly disciplines
 - Hardware for AI: A team of NIST researchers, mostly physicists and mathematicians, are working on fabricating and measuring new brain-inspired circuits and architectures based on novel devices to enable future generations of AI.
- U.S. Government AI Standards Coordinator
 - Outreach to connect with all known federal efforts relating to AI standards development and use with the goal of community participants leveraging and learning from the successes of other participants.
 - In collaboration with OSTP and the National Science and Technology Council (NSTC) Subcommittee on Machine Learning and AI (MLAI), plan and execute an international campaign coordinating effort on AI standards development

AI Provisions in the National Defense Authorization Act of 2021:

- There are numerous provisions and consequences related to AI that were in the legislation, which is approximately 4,500 pages. The most important provision is the creation of a new National AI Initiative Office led by the White House to coordinate federal support for AI research and development, education and training, research infrastructure, and international engagement.
- Codifies the NSF AI Institute. It includes provisions that established a National AI Research Resource task force and formalized the National AI Research Institute effort.
- NIST is given 2 years to develop an AI Risk Management Framework

National AI Initiative Act (Title LI, Sec. 5101):

- Establishes an office within the White House to serve as a point of contact for federal AI activities as well as public and private entities that may be involved. The office is comprised of two core organizations: Interagency AI Committee and the National AI Advisory Committee. The committee has 2 years to develop a strategic plan. There was an AI Select Committee, which is being re-chartered to the Coordination by Interagency AI Committee. The National AI Advisory Committee will be established by the Secretary of Commerce in consultation with others.
- NIST Activities:

- Title LIII, Sec. 5301: NIST is to expand its mission to include advancing collaborative frameworks, standards, guidelines for AI, supporting the development of a risk-mitigation framework for AI systems, and supporting the development of technical standards and guidelines to promote trustworthy AI systems.
- Develop a voluntary risk management framework for trustworthy AI systems and regularly update it.
- Participate in standard setting organizations.
- Develop data-sharing best practices and best practices for documentation of data sets.
- Conduct stakeholder outreach, which NIST is already doing.
- Other AI provisions:
 - National Academies AI Impact Study on Workforce
 - National AI Research Resource Task Force: How should resources be established and sustained? Develop a plan for a national research Cloud that would allow access to the data.
 - National AI Research Institute: They are codifying the AI Research Institute that NSF started. Each institute will be funded for up to 5 years, with \$4 million per year.

National Security Commission on AI: Final Report (March 1, 2021)

- 750+ pages, with 16 chapters in the main report, providing topline conclusions and recommendations. It is accompanied by a blueprint for action, outlining detailed steps that the U.S. government should take to implement the recommendations.
- Chapter 7: Establishing Justified Confidence in AI Systems
 - Robust and Reliable AI
 - Human-AI Interaction and Teaming
 - Accountability and Governance
 - Testing and evaluation, verification, and validation: This includes a specific recommendation for NIST to provide a set of standards, performance metrics, and tools for qualified confidence in AI models and data training.
 - Leadership
- Chapter 8: Upholding Democratic Values: Privacy, Civil Liberties, and Civil Rights in Uses of AI for National Security
 - Invest in and adopt AI tools to enhance oversight and auditing in support of privacy and civil liberties.
 - Improve public transparency about how the government uses AI.
 - Develop and test systems with the goal of advancing privacy preservation and fairness.
 - Strengthen the ability of those impacted by government actions involving AI to seek redress and have due process.
 - Strengthen oversight mechanisms to address current and evolving concerns.
- Recommendations for NIST: NIST is mentioned quite a bit in the report but at a very high level.
 - Continue to support the development of best practices for data, model and system documentation.
 - Provide a set of standards, performance metrics, and tools for qualified confidence in AI models, data, training environments, and predicted outcomes.
 - Facilitate third-party test centers for AI systems: This recommends a role for the National Voluntary Laboratory Accreditation Program (NVLAP) to facilitate test centers for AI systems.

Ms. Tabassi mentioned that she had asked one of the commissioners from the National Security Commission on AI to brief the ISPAB, but they were busy. She may ask again for a future meeting.

Discussion:

- The Chair noted that this is a fast-moving area with a lot of progress and technical development taking place in industry, while at the same time we are seeking principles or best practices for achieving the building blocks. It seems like a risk area that either we get requirements that aren't achievable or take us in the wrong direction or we get systems fielded that ignore important considerations. How do you avoid these outcomes?

Ms. Tabassi said that in a perfect world you want to have very solid scientific foundations, but the industry is moving fast. There is the risk of adopting something that is less than optimal, but that is a lower risk than the risk of adopting something that's completely wrong. This problem almost always exists, but it is exacerbated with AI because of the pace of the technology. There are things we can do like create a checklist and start documenting what's happening from the design to implementations to development and deployment. NIST is looking for guidance and wisdom from the board. Produce as many documents as possible and keep them regularly updated as the technology changes.

- Ms. Fitzgerald-McKay asked how we assess U.S. leadership in AI and what the risks are of not taking a leadership role.

Ms. Tabassi said that question can be addressed in many ways. One way is making sure that standards are developed in a consensus manner and are globally relevant. Another way is in the area of policy and regulations. The EU released an AI White Paper in February. In April they are coming up with another set of documents. They are going toward identifying some sectors as high risk and in need of regulation. This can have an impact on U.S. industries. Additionally, Ms. Tabassi is a big believer in the U.S. research eco-system, which is strong. There have been some very positive steps in research. NeurIPS, a very respectable academic conference for machine learning, started a new requirement about understanding the impact of research, whether it is positive or negative.

- Ms. Hallawell asked about implications of AI for cybersecurity. Should there be any branching off in NIST's AI work to think about it specifically for security? A lot of enterprises are marketing tools that say they do AI. Is there a buying guide for federal agencies or civilian organizations? Any criteria to know if it is even AI?

Ms. Tabassi said the Consumer Product Safety Commission is working on a way to educate consumers. NIST is not working on that right now. An OMB regulatory memo from November 2020 tried to set forth some of those guidelines. An executive order in December set out high-level principles – being lawful, purposeful, accurate, reliable, private, etc. Both of those are still high-level, value-based statements.

DOD Cybersecurity Maturity Model Certification (CMMC)

Stacy Bostjanick, Acting Director, Supply Chain Risk Management, Office of the Under Secretary of Defense

Mr. Scholl introduced Stacy Bostjanick, Acting Director of Supply Chain Risk Management at the Office of the Under Secretary of Defense, who presented some background on CMMC:

- An IG investigation and a Navy cyber-readiness review looked at compliance with NIST SP 800-171 and found that companies were not doing what they said they were doing. Either they really didn't understand what they were complying with, or they were just checking a box and not really paying attention. The Secretary of Defense at the time decided that compliance had to be validated.
- In August 2019, they published the first model revision 4 for comment. They went across the world and gathered all the cyber standards and requirements they could find and included them. They received about 2,500 comments from the public.
- On January 30, 2020, they published the final model. They did not want to require anything that was so expensive or difficult that it would discourage a company from participating in DOD contracts. The model ended up being for CMMC Level 3, which is what they considered the basic standards for handling controlled unclassified information (CUI). It was closely aligned with the NIST SP 800-171 requirement. They added 20 more requirements to NIST SP 800-171 for CMMC Level 3. For CMMC Levels 4 and 5, which would have heightened security, the model is closely aligned to NIST SP 800-172.
- A lot of the assessment and verification work was being done by the DIBCAC [Defense Industrial Base Cybersecurity Assessment Center] team and DCMA [Defense Contract Management Agency]. They recognized that it wasn't scalable with 300,000+ companies, so they issued an RFI for companies to help manage the CMMC effort. They received 36 responses, but no single respondent could cover the full capability.
- In November 2019, they had asked industry to come together to form an accreditation body (AB) to help manage the assessments and verifications. That became the CMMC AB. They worked predominantly on a voluntary basis and then moved to a no-cost contract. They are not being paid for any of their work and have worked thousands of hours to set up a framework.
- They took a cue from FedRAMP in terms of third-party assessment organizations. They met with FedRAMP to find out what worked well and what didn't.
- They also recognized the need for a DFAR [Defense Federal Acquisition Regulation] rule to be able to require this as an enterprise requirement across the DOD. They started with a proposed rule. OMB recognized the national security interest in CMMC and made it an interim rule, which essentially allowed them to put it into effect prior to going through the entire process of finalizing and formalizing it.
- On November 30, 2020, it became an interim rule, allowing them to move into a pilot phase to require CMMC as a condition of award in select pilot programs. Prior to that, they also engaged in pathfinder activities as a risk reduction with the Missile Defense Agency. They came up with a mock scenario for a contract and a program. They worked closely on four tabletop exercises and some mock assessments. There were two teams – the contractor team and the government team. They came up with model language for an RFI and allowed the prime contractor to provide feedback. They crafted model RFP language and then they moved to a post-award conference. They also worked through a dispute resolution scenario.
- The CMMC program management office provided the CMMC AB with learning objectives, and the AB turned that into a curriculum and a course. They trained the DIBCAC assessors that are currently doing assessments for DOD.

-
- They looked at a company that had a DIBCAC high assessment to see what it would look like if a CMMC assessor conducted an assessment. They did a full Level 3 mock assessment and two Level 1 mock assessments. Their expectation was that CMMC Level 3 companies are only about 20 percent of the DIB. The rest will only ever have to be CMMC Level 1 compliant. One company failed a Level 1 mock assessment but only because of small, inexpensive things. They forgot to put a log at the front door, which meant they didn't have positive access control to their space, and they left a brick in the back door for people to go outside and smoke.
 - They had agreed to a 5-year, phased-in approach, with up to 15 acquisitions the first year and 479 acquisitions by the fourth year. By FY 2026, when it goes into full force, the model will apply to all contracts.
 - Certifications will be valid for 3 years. However, if a company is in the middle of a certification and is going to be awarded a contract that is considered particularly sensitive, the DCMA DIBCAC team would be able to do a spot check.
 - For C3PAOs, they recognize that the data and information assessors will access is sensitive. They have required that all C3PAOs go through CMMC Level 3 assessments themselves, which will be performed by DIBCAC assessors.
 - Because the 7012 clause is not going to go away, it will still be included in contracts. If a company experiences a breach, they will have to report to DC3 and the federal government what occurred and how they were impacted. Then the DMCA DIBCAC team will triage it because they will want to know whether 1) the assessor did something wrong, or 2) the company became negligent after assessment, or 3) the model did not address that particular risk.
 - They will review the model at least annually as the risk changes.
 - In the SolarWinds scenario, CMMC probably would have given a someone at Level 3 the ability to see the activity. At Level 4 and Level 5, they might have had measures to prevent SolarWinds.
 - CMMC is considered foundational: They do not want to be able to negotiate away security.

Discussion:

- The Chair said that when they heard about CMMC in 2019, one concern was that there needed to be a degree of piloting to make sure it worked before being rolled out to the world. Pilots should 1) test the assessment and 2) test the security of the systems to see if CMMC actually was correlated with improved security and improved resistance to attack. Conducting a root cause analysis would involve looking at incidents and seeing why they happened and what in the process or tools would fix it. How far down into implementation do CMMC controls go? Would assessors detect the brick in the back door?

Ms. Bostjanick said yes. The process allowed them to tweak the training to make sure it includes real-world experiences. If you need to be CMMC Level 3 compliant, what do you need the assessor to look at? They will define the boundaries. They might say this one portion of the business is where the CUI is handled. They can decide what they have looked at. With Level 4 and Level 5, most companies will have a cut-out enclave to handle that level of sensitive information, but the assessors will be trained to look at all of that.

There are ongoing conversations with industry to strike a balance between protection and not being so onerous as to drive a company out of business or drive them crazy. We must be doing [the CMMC] right because people on both sides are complaining. It is a crawl, walk, run

approach. They started the 7012 clause effort around 2012-13 and didn't get something in place until 2017. From 2013 to 2020, companies started to see they were being hacked.

CMMC is a foundational portion of a larger effort of supply chain risk management. The hope is that many companies say that they're paying attention, look at the assessment guides, recognize where their gaps are, and start taking steps to close them because they understand it's the smart thing to do, not because they're required to. It does no good to develop intellectual property and use your innovative capabilities if adversaries steal it.

- Mr. Venables asked how much of the inability of companies to implement cybersecurity is a result of bigger IT problems because they haven't upgraded their IT for decades. Are we seeing that in the defense contractor space where there is a need to modernize IT as a foundation for implementing cybersecurity? Or is it more complex?

Ms. Bostjanick said she hasn't heard that complaint. However, small, innovative companies say they have their nephews come in and update their IT. For COTs products, they're not requiring this. CMMC Level 1 is the basic hygiene. One of the challenges the department is facing is in training its own people and identifying the CUI. One of the problems is that program managers and the primes are not taking the time to disaggregate the data and limit it to what the subcontractors need. There are people who are struggling because they haven't upgraded their IT and people struggling because they don't even have much of an IT infrastructure in their company.

- The Chair asked about the SolarWinds incident and why a Level 3 would have been able to see the activity and why Level 4 and Level 5 would have been able to stop it.

Ms. Bostjanick said she is not a technical guru but her team put together a briefing on this. Level 3 protections would have allowed them to see the movement of the data on the network. Level 4 and 5 controls – 24-hour SOC and others – could have possibly stopped the activity. State actors, like China and Russia, do have heightened capabilities.

- The Chair noted that the speakers from the University of Maryland talked about compliance programs and the challenges of getting from requirements to actual security impact. One of the failure modes we see a lot is that people go through the requirements one by one but don't achieve an integrated secure system. In another, because compliance requirements are individual sentences in the passive voice, companies wind up being able to comply despite having a system that's still vulnerable. Have we gotten past that with CMMC?

Ms. Bostjanick said SP 800-171 is a subset of SP 800-53. When they tried to impose SP 800-53 on the DIB originally, they said it was too difficult to meet. They tried to provide a holistic view in the assessment guides for each level. The Level 1 and Level 2 assessment guides have been written at an eighth-grade level. Level 3 gets more complex. They worked with the CMMC AB to implement this, and they are doing training for the assessors, the C3PAOs, Registered Practitioners (RPs), and Registered Providers Organizations (RPOs).

You are not allowed to be a consultant and an assessor at the same time. They have provided training to make sure companies understand CMMC and the methodology and intent. Ms. Bostjanick said she has received negative feedback on LinkedIn because she suggested that if you are a company looking for a consultant or a C3PAO for your assessment, you must choose from the CMMC AB marketplace because they will only accept certification from C3PAOs that the CMMC AB has trained and accredited.

Companies have said they've paid thousands of dollars to consultants who did nothing. Buyer beware. They still struggle with companies that are focused on the checklist. But when they have an assessor come in, the assessor will give them guidance on why they failed and what they need to do to rectify it. There will be a couple frustrated companies that are going to rush to certification without thinking it through.

- The Chair said that the root cause analysis is very important. Don't look for success – keep looking at failures and try to prevent them from happening again.

Ms. Bostjanick agreed and noted that the model is not going to stay static. It will be a mechanism to help protect, but it's not going to be the end all and be all. They are heartened that a couple primes that had put most of the CMMC requirements into effect were somewhat protected and were able to shut it down when they saw the activity related to SolarWinds. The hope is that companies move in this direction because they realize they have to for their business base. Not losing data can give them an advantage in the market. We are hoping to come up with tools that will reduce costs. Also, Trusted Capital will help make sure there are good investments in the industrial base and not just adversaries investing in the industrial base.

The Chair thanked the speaker and recessed the meeting for a 6-minute break.

Cybersecurity IOT Act of 2020 and NIST Efforts

Katerina Megas, Program Manager, Cybersecurity for IoT Program, NIST

Kim Schaffer, IT Specialist in the Security Components and Mechanisms Group, NIST

The Chair welcomed Kat Megas and Kim Schaffer of NIST.

Ms. Megas presented an overview of NIST's work applicable to IoT cybersecurity, noting that for years much of the work they have been doing is broadly applicable in support of IoT cybersecurity.

- Program Principles Guiding Efforts: They established these principles about 4 years ago when the program was started.
 - Risk-based understanding: Focus on how IoT characteristics affect system and organizational cybersecurity risk.
 - No one-size-fits-all: Allow for diversity of approaches and solutions across industries, verticals, and use cases.
 - Ecosystem of things: No device exists in a vacuum, so look at the entire ecosystem, not just endpoints.
 - Stakeholder engagement: Collaborate with diverse stakeholders regarding tools, guidance, standards, and resources. They started out nearly 4 years ago by hosting an IoT security colloquium.
 - Outcome-based approach: Specify desired outcomes and allow providers and customers to choose best solutions for their devices and environments.
- Key Events:
 - NISTIR 8201 (December 2017)
 - NISTIR 8228 (June 2019)IoT Cybersecurity program's first publication. They wanted to help federal agencies understand how to make use of existing guidance.
 - Focuses on what is different about managing risks associated with the use of IoT. Where might there be challenges or assumptions that a federal agency might make about the capabilities of the device?

- Frames IoT risks and challenges in the context of implementing SP 800-53 controls and the Cybersecurity Framework. Much of NIST guidance is developed as technology agnostic and applicable to managing security risks associated with the use of IoT devices.
- Customers dependent on security capabilities of IoT devices: There is still an inherent dependency between an organization being able to achieve its security goals and the devices they were integrating into the high-level enterprise systems.
- NISTIR 8259/8259A (May 2020)
 - Three public workshops, two public comment periods, and over 600 comments
 - Cybersecurity recommendations for IoT device manufacturers
 - Activities for manufacturers to incorporate into product development lifecycle
 - Six core cybersecurity capabilities for IoT devices
- Federal Profile Workshop (July 2020)
 - Published on GitHub analysis of SP 800-53 controls dependencies on IoT device capabilities. Suggested this be a ‘catalogue’ for agency use/
 - Takeaways: Confirmed device-centric approach was useful, non-technical dependencies need to be identified, and confidence mechanisms are desired for the market but more discussions is required.
- Four Public Drafts (December 2020)
 - Non-Technical Supporting Activities Baseline recommended for all IoT device manufacturers
 - The process NIST followed to adapt the baseline to federal agency use case
 - Starting point for agencies in a federal profile identifying the key capabilities likely needed to support agency implementation of low baseline
 - Guidance for Federal Agencies with considerations for IoT risk in agency RMF [Risk Management Framework] processes and how to develop requirements for IoT devices leveraging catalogue and federal profile
- IoT Cybersecurity Improvement Act (December 2020): Directs NIST to come up with public standards and guidelines for the federal government on the appropriate use and management by agencies of IoT devices.
- IoT device cybersecurity should be addressed within a risk management hierarchy from enterprise level through organization, system, and finally component level, where IoT devices are understood as system components with a distinctive set of risk characteristics.
- Working description of IoT (June 2020)
 - NISTIR 8259: Process for manufacturers.
 - They did not set out to create a definition for IoT or IoT devices. Most of the feedback from industry was that NIST does not need to define IoT to be able to talk about the cybersecurity risks and concerns. But for the purposes of documentation, they came up with a high-level “description of IoT devices,” which is intended to be able to be applied widely. The IoT Cybersecurity Improvement Act borrowed the definition from NISTIR 8259.
 - Described IoT devices as having at least one transducer for interacting directly with the physical world *and* at least one network interface for interfacing with the digital world.
- NISTIR 8259A (May 2020)
 - Technical baseline (device identification; device configuration; data protection; logical access to interfaces; software update; cybersecurity state awareness)
 - They set out to identify the core baseline, understanding that one size does not fit all. Early on, they talked about the need in the future to likely adapt the core baseline or extend it. They didn’t think it was up to them to define profiles.

-
- They published recommendations which can be used across a wide range of IoT devices.. Profiles can be developed building on the core baseline to define the market- or vertical-specific needs.
 - Four additional draft publications were released to create a framework for profiling requirements for devices (December 2020):
 - NISTIR 8259B: Identified non-technical capabilities that might be broadly applicable and could be considered core. Recommendations cover training and awareness; disposal practices; physical protections, reporting and much more. Non-technical core baseline: documentation; information and query reception; information dissemination; education and awareness.
 - NISTIR 8259C: They documented the process used to develop the federal government profile for the baseline. If there were other organizations interested in developing similar guidance, they could look at this.
 - NISTIR 8259D: Profiles and adapts the core baseline in 8259B to federal agency needs. They started out with NIST guidance that was not cyber specific for primary source documents. (Cybersecurity Framework; NIST SP 800-53 Rev 5; low impact baseline from NIST SP 800-53B; technical capabilities from NISTIR 8259A; non-technical capabilities from NISTIR 8259B) They identified an additional technical capability for IoT devices, meaning that the IoT device can operate securely by protecting its hardware and software integrity and security utilizing system resources, managing communications.
 - SP 800-213: IoT Device Cybersecurity Guidance for the Federal Government: Provides guidance for federal agencies to consider as they establish requirements. This includes an IoT Device Security Capabilities Catalogue, which provides a taxonomy of requirements they can draw from when the agency determines the baseline is not sufficient. The baseline is a subset of the catalogue.
 - First OLIR [Online Informative References] mapping of NIST recommendations to standard: This allows them to have informative references that have a relationship to the NIST guidance. It will be an important tool as NIST puts out various guidelines.
 - They have started engaging with some NCCoE projects, beginning with the Secure Telehealth Remote Patient Monitoring Project. They want to map the recommendations they made in 8259A.
 - Key areas of expanding work
 - Confidence mechanisms for the marketplace: White paper on how to get to confidence in the security of IoT devices
 - Consumer devices apply the guidance in NISTIR 8259 (Updates to NISTIR 8267 Security Survey of Consumer Home IoT Products; Workshop on Cybersecurity Risks in Consumer Home IoT Products, October 2020)
 - Preliminary high-level themes in public comments:
 - Various views on how the risk of adding an IoT device to a government network should be characterized
 - Significant concerns about understanding fragmentation, including market fragmentation, policy fragmentation, and different agencies defining IoT cybersecurity requirements differently. They provide a catalog with some consistency in language around requirements, but they do not specify what requirements are applicable in which scenarios. Some commenters are concerned that they may end up with a significantly fragmented set of requirements from different agencies.
 - Many IoT devices are too constrained to support the requirements (precluding the use of large number of IoT devices by government)

- Templates for requirements for different types of devices are needed.
- Call to make distinctions among device “types”
- Tentative public workshop: April 2021

Mr. Schaeffer presented an overview of the Cybersecurity IoT Act of 2020 and NIST’s role:

- Section 5 of the Act addresses guidelines on the disclosure process for security vulnerabilities relating to information systems, including IoT devices. NIST was given 190 days – by June 2, 2021 – to develop and publish guidelines for the reporting, coordinating, publishing, and receiving of information about security vulnerabilities. The guidelines shall, to the maximum extent practicable, be aligned with industry best practices and Standards 29147 and 30111 of the International Standards Organization (or any successor standard) or any other appropriate, relevant, and widely used standard. The guidelines shall include example content on the information items.
- There are many different levels that have to be addressed for this to work. It is likely to be complicated. Mostly, it is about how to make sure your reports go to the right people – that your processes are correct for each individual team of developers to find, address and incorporate fixes through the systems that use that software.
- The plan is to “publish” a guideline based on ISO 29147 and 30111 as a normative reference with tailoring for U.S. government specifics initially from OMB Policy 20-32, BOD 20-01. Questions include: What does this mean for access to ISO documents? What gets tailored?
- To the maximum extent practical, they would like to use whatever guidelines are in place. The problem gets to be with the example content. What does it mean? If they have to develop a lot of content and it has to be published at the right level and completed in 6 months, it is a hard effort. Questions include: Is the current CVE Numbering Authorities participation sufficient? Are the relevant standards effective for these purposes? Is content data level or descriptive of type?
- A person reporting could send the report to a CNA, although it could also be sent to a supply chain vendor or another government office. It could possibly be reported directly to the software vendor developers themselves. Those are three ways it could go. They are trying to make sure the report goes to the right person as quickly as possible, but some reporting needs to be done so there is some overall management. There are basically two coordinators – one at the federal level (basically a list maker and taker and source of knowledge of the status of those things) and one in the software program office who is responsible for trying to address the particular issue.
- References for Use and Going Beyond the June 2 Requirements:
 - Workshop/Conference for public discussion and inputs: They have already started discussions, especially with DHS and DOD, about processes and how the individual software shops work – in a software development office, at the agency level, or in a small, one-off type custom software.
 - IR 8246: Collaborative Vulnerability Metadata Acceptance Process (CVMAP) for (CNAs) and Authorized Data Publishers
 - OSS Vulnerability Guide <https://github.com/google/oss-vulnerability-guide/blob/main/guide.md>: This deals with how open source software is trying to address this same vulnerability disclosure.
 - SP 800-61, Rev. 3 (underway): Computer Security Incident Handling Guide
 - Draft IR 8138 -Vulnerability Description Ontology (VDO): A Framework for Characterizing Vulnerabilities
 - IR 8011, Vol. 4 -Automation Support for Security Control Assessments: Software Vulnerability Management
 - SP 800-40, Rev. 3: Guide to Enterprise Patch Management Technologies

- Forum of Incident Response and Security Teams: Vulnerability Reporting and Data eXchange SIG (VDRX-SIG) CNAs; Common Vulnerability Scoring System
- Others: SWIDS; SBOM etc.
- Existing Industry and Agency Programs and Policies to reflect
- Bringing it all back to ISO or other SDOs
- Potential R&D Challenges:
 - Do we have the right data in the right formats for the right purposes?
 - Are we automating in an interoperable format for integration with our other essential response capabilities?
 - What measures/metrics can be understood from the generated information? How can we use this to improve vulnerability identification, incident response, and recovery?
 - Have we correctly identified incentives for participation and reduced barriers for submissions, not just technical risks but business, customer, legal, and economic risk?
 - What is the needed adoption/use to achieve “effective mass?”
 - What are different measures?
- Even if this gets done by June, it probably won’t be finished. There will be an awareness campaign rolled out to make individual software vendors – internal vendors – understand that they have the capability to process reports from multiple areas – to try and coordinate the multiple vulnerabilities, if possible.

Discussion

- The Chair noted that the intent of the coordinator role as specified in ISO 29147 and 30111 is as an intermediary. If a vulnerability researcher finds a problem somewhere, rather than report it directly to the developer or vendor, he or she can choose to work through an intermediary that will do the reporting. CERT CC is a historic organization that filled that coordinator role. Vendors reporting to other vendors may be a supply chain relationship, but it may also be a “multi-vendor vulnerability disclosure.” With the Specter and Meltdown vulnerabilities that affected processor chips and the software that ran on them, every operating system supplier had to coordinate on addressing and remediating a common class of vulnerabilities. The basic model is simple: Find vendor, report vulnerability to vendor, vendor fixes it and releases a patch. But the reality can get gnarly.

Mr. Shaeffer said it is a huge problem and a huge education for some of the smaller internal development. Many are still protecting valuable information. They need to be made aware at a minimum. Chances are NIST probably won’t be dealing with a lot of hardware developed directly by the government. It takes a long time to get this process incorporated into development.

The Chair said the primary scenario to think about is getting a basic capability out soon for software the government develops, operates, or owns and exposes to the public.

- Ms. Hallawell noted the focus in the legislation on the device manufacturer making the device securely. Is there anything else that talks about the responsibility of an agency or downstream enterprise to use the device securely?

Ms. Megas said SP 800-213 is directed at federal agencies, reinforcing all the things an agency needs to be doing and trying to position them to acquire devices that can support what they need to be doing. Security doesn’t reside completely on the device or the device manufacturer.

The Chair thanked the speakers and wished them luck in their challenging set of tasks. He recessed the meeting for a lunch break.

Facial Recognition Testing, Accuracy and Bias

Patrick Grother, Computer Scientist, NIST

Craig Watson, Image Group Manager for Biometric Research, NIST

The Chair introduced Patrick Grother and Craig Watson of NIST.

Mr. Grother began the presentation with an overview of Face Recognition Technology:

- Face recognition works by comparing faces. It does not have an in-built memory of anyone in particular. It renders a similarity measure of two photographs and helps determine if they belong to the same identity. It is a one-to-one process, involving typical hypothesis testing, similar to radiography or many battlefield detection tasks.
 - False positives and false negatives are part of the game.
 - Human capability to determine the same or different identity when looking at pairs of images is a non-trivial task. The leading automated face recognition algorithms today will get 100% correct classification, which is quite rare for humans.
 - With automated one-to-one face recognition technology, both images go into an identical feature extractor, based on neural networks, which extracts vectors of numbers that are representative of identity. If the vectors are similar, there is a high similarity score for the two images. The technology is built on discrete neural networks, which are built on machine learning and AI technologies. They are not commoditized, and they embed considerable trade secrets. In principle, they can leak privacy information.
 - There is a bigger marketplace built on a one-to-many capability, which compares a photo to a database that's been pre-enrolled. It will do one-to-one comparisons and produce a list of similarity scores coming out with an identity. In principle, it can make a false negative and a false positive.
 - Face recognition has undergone a revolution of sorts in the last decade based on a new generation of convolutional neural networks (CNN), resulting in longitudinal gains.
 - False negatives have reduced significantly in the last few years. There is steady improvement in accuracy by training neural networks to be more tolerant of image quality-related defects and more tolerant of aging effects.
 - Cross-pose matching is the Holy Grail of face recognition research, and the ability to recognizes people off angle has increased.
 - Enablers of better face recognition include the ability to scrape social media, CNNs, digital cameras, fast and robust machine learning tools, publication of databases of images for developers, continued development of CNN architectures, and more.
 - Deep Convolutional Neural Network (DCNN): Industry these days has been quite candid in indicating that the technology is built on convolutional neural networks. A CNN accepts input typically at image width by height, number of color channels, which goes into a composed function. Some algorithms are open source, including one from Imperial College London, which was released publicly 2 years ago. It was very competitive with leading industrial players at the time.
- Face Recognition Vendor Testing
 - Face recognition vendor testing (FRVT) started in 2000 and was put into a continuous mode in 2017. Any developer could send NIST an algorithm for evaluation – no more frequently than every 4 months though. It involves a series of benchmarks that are public, independent, open worldwide, free, uniform API, archived sequestered images. They run the algorithm on an archive of sequestered images, which are not available to the developers. Other kinds of

- testing also exist. Examples of images used in tests include 1) mug shots and images collected on the southern border and 2) other databases they can't reveal publicly.
- Four ongoing benchmarks:
 - o Verification
 - o Search Performance
 - o Morphed Photo Detection
 - o Automated Quality Assessment
 - Reports: From the four benchmarks, they have produced regular updates to six different reports.
 - o Performance of 1:1 Verification Algorithms
 - o Performance of 1:N Identification Algorithms
 - o Demographic Effects in Face Recognition
 - o Performance of Morph Detection Algorithms
 - o Performance of Image Quality Assessment Algorithms
 - o Performance of Face Recognition with Face Mask
 - o Upcoming: Performance of face recognition on twins
 - o Upcoming: Use of face recognition in paperless travel
 - They are not working on topics of presentation, attack detection, or on privacy-related topics such as de-identification of images and evaluation of protected templates.
 - Performance Tables
 - There is a wide range of accuracy across the industry.
 - They produce about a dozen tables.
 - Twins, the Forgotten Demographic: An unsolved problem in face recognition.
 - Twins are important because face recognition algorithms should not have false positives on twins. Twins are quite common and are becoming more common, particularly non-identical twins. What happens when you enroll a database of 1.6 million identities and you put one of the twins in the database?
 - They used photos from Notre Dame collected for the Twins Day Festival.
 - CDC maintains good statistics on the number of twins in the population.
 - Identical twins: A rank 1 hit with a high score tends to come back: In principle, that could cause trouble for the twins in a criminal investigation, which you would not expect to happen when comparing fingerprints.
 - Fraternal twins: The scores are lower. At least one algorithm does not identify the twin.
 - Trade-off between false negative and false positive rates: If you try to run a face search on all the individuals in a small town, for example, you would not be able to achieve a low false positive rate. The same would be true with a national search. Twins manifest as an inability to get to very low false positive rates using current technology.
 - Demographic Effects in Face Recognition
 - NISTIR 8280 (December 2019)
 - A highly influential report had been published in 2016 by the Georgetown Law School Center on Privacy and Technology, "The Perpetual Line-Up." It addressed the use of face recognition in criminal justice. They found evidence that face recognition would be biased against certain demographics. It was followed up by a study from MIT, which did not actually look at face recognition technology, but instead at face classification technology.
 - Publicizing of high error rates on dark-skinned females in the processing of single images (not face recognition comparing images).
 - Those were driving factors for an effort to characterize face recognition performance on different demographic groups.

- December 2019 NIST demographics report
 - o They used 187 algorithms from 99 developers, mostly commercial. They applied it in one-to-one mode and one-to-many mode using different data sets.
 - o They tried to describe that different demographic factors would have different relevance to different applications. They used four different databases and were able to look at demographics characterized by sex, age, and race.
 - o There is a prominent example of a declared instance of demographic differential, which comes from Apple's implementation of face authentication on their iPhone X. Apple, to its credit, pointed out that there is a limitation in the technology for twins, siblings, and children under the age of 13
 - o False negative differentials are low because face recognition accuracy is very good. So the differentials between demographics are low.
 - o False positives rates vary widely between different demographic groups. False positives are the contemporary problem for face recognition.
- False Match Rates
 - o Across the board, false match rates vary depending on age.
 - o For white males hailing from Eastern Europe, the false match rate is 1 in 30,000.
 - o For Chinese women over the age of 65, the false match rate can be as high as 1 in 30.
 - o They published the data and implicitly challenged the community to try to address the differentials, and there have been some commercial efforts to do so since 2019.
 - o They categorized 24 countries into seven regions of the world. They found significant variation in false positive rates among regions.
 - o They found elevated false positive rates in West Africa, East Africa, and East Asia. It suggests that the training data used for development of face recognition algorithms and the location of the developers impacts these outcomes, but this is not formally proven. It suggests that if you are able to use diverse data, then in principle you could address the false positive issue.
- They made a key point that any demographic differential may or may not be material, depending on the wider range of applications that the algorithms are being used in. You have to think through what the errors could mean in practice. Applications include dispensing drugs, paperless boarding, and watchlist.
- Summary
 - Leading contemporary algorithms are very accurate and increasingly tolerate poor quality.
 - However, they distribute errors inequitably across demographics. False positive differentials are much larger than false negatives, particular for women, Asians and Africans (with some exceptions), and the old and very young.
 - Algorithm matters: Accuracy varies, and demographic sensitivity varies.
 - Application matters: Errors have to be thought through.
 - There has been incomplete reporting in the press and academia
 - Since 2019, some developers have attempted to address the differentials.
 - NIST recently developed a summary indicator of differentials, but it is not yet published.

Discussion:

- Mr. Groman asked about the difference between one-to-one and one-to-many algorithms, and why it's relevant.

Mr. Grother said one-to-many algorithms are controversial for civil liberties reasons, but they are capable of going through extraordinarily large databases with very little increase in error rate.

- Mr. Maughan asked about the acquisition of face recognition technology. Do organizations like DHS, DCP, state and local law enforcement just buy from a vendor, or is NIST involved in the acquisition decisions?

Mr. Grother said all of the U.S. government places contracts for procurement, and some look at NIST test reports. NIST talks to people across U.S. government but never gets involved in procurement.

- Ms. Fitzgerald-McKay asked if tests are available for those involved in procurement.

Mr. Grother said the NIST disposition is to publish everything, so they release big reports, which are open worldwide.

- Mr. Groman asked about the one-to-one scenario involving a passport photo presented at the airport, which is then compared to a photo taken at the terminal. In that scenario, there is also the possibility for a human to check in the event there is no match – an internal safeguard. What other scenarios are there for that?

Mr. Grother said there are a number of one-to-one applications. U.S. government ID cards include a face image, which could support one-to-one authentication for logical access or physical access. Another example is non-repudiation, such as a pharmacist dispensing drugs who needs to be able to sustain a claim that he wasn't the one who gave someone opiates. Typically, they use a fingerprint for that, but they could use a face image.

DHS and DOT efforts for PNT in EO 13905, Strengthening National Resilience through Responsible Use of Position Navigation and Timing Services

James Platt, Chief for Strategic Defense Initiatives, DHS

Karen Van Dyke, Director of the Office of Positioning, Navigation, and Timing, DOT

The Chair introduced James Platt, Chief for Strategic Defense Initiatives at DHS, and Karen Van Dyke, Director of the Office of Positioning, Navigation, and Timing at DOT

Ms. Van Dyke opened the presentation with an overview of GPS and some of the current challenges and threats.

- The Department of Transportation (DOT) has been the civil lead for the Global Positioning System (GPS) for several decades. In next-generation aviation, positive train control, intelligent transportation systems, and the maritime environment, we need trusted sources of navigation and timing. It is critical to reduce delays, prevent crashes, and increase efficiency, among other things.
- GPS/GNSS [Global Navigation Satellite System] challenged environments:
 - Ionospheric disturbances
 - High accuracy with integrity
 - Timely notification of misleading information
 - Underground/indoors
 - Urban canyons
 - Inaccurate and out-of-date maps, are a growing problem, including trucks crashing into bridges and underpasses. It isn't the GPS signal that is causing the issue; it's the interpretation of the information being provided to the user.
- Existing GPS/GNSS Threats: They have been increasingly concerned about attempts to disrupt or manipulate the GPS signal.

- Jamming is intentionally produced RF waveforms that have the same effect as interference.
- Spoofing is a bigger threat than jamming. If you are providing false information to the user and they are not able to detect that, it is extremely hazardous. It can deny, degrade, disrupt, or deceive a receiver's operation and can have a range of effects, from incorrect outputs of positioning, navigation, and timing to receiver malfunction.
- Executive Order 13905 (February 12, 2020)
 - This was long needed to focus on strengthening national resilience through responsible use of PNT services. We need to understand the end user requirements and whether GPS is the best solution to meet those requirements given the vulnerabilities.
 - DOT is not in any way backing away from use of GPS. They refer to it as the cornerstone of PNT architecture. However, they need to look at other technologies and how to drive down the level of risk.
 - Purpose: Foster responsible use of PNT services by critical infrastructure owners and operators to strengthen national resilience
 - Policy: Ensure that disruption or manipulation of PNT service does not undermine the reliability or efficiency of critical infrastructure; raise awareness of the extent to which critical infrastructure depends on PNT services; ensure critical infrastructure can withstand disruption of manipulation of PNT services; engage public and private sectors to promote responsible use of PNT services
 - Key Actions
 - PNT Profile development: NISTIR 8323: Foundational PNT Profile
 - National R&D Plan on PNT Resilience, released by OSTP
 - Pilot Programs/Vulnerability Assessment/Testing
 - GNSS Independent Source of UTC
- DOT Pilot Program Overview
 - The executive order requires pilot programs conducted by the DHS, DOT, and the Department of Energy. DOT is focused on the maritime environment, where there have been increasing occurrences of GPS jamming and spoofing, not so much in the United States but particularly in Russia and China.
 - DOT partnered with the Maritime Administration to better understand the dependencies as well as the impacts when a signal is jammed or spoofed.
 - Conduct stakeholder engagement
 - Evaluate complementation PNT technologies suitable for the maritime environment
 - Develop a jamming and spoofing detection capability
 - Results will provide insights to support the development of PNT profiles for maritime applications, as well as to inform additional PNT R&D.
 - Results and lessons learned may benefit PNT resiliency for other modes of transportation (aviation, rail, vehicles, and pipeline).
- GPS Jamming and Spoofing in the Maritime Environment
 - DOT public workshop (December 3, 2020: More than 400 people registered, including ship captains who had experienced jamming and spoofing. There were briefers from the NSC, OST-R, MARAD, USCG, and the Volpe Center. Briefers from industry included Maersk, APL Maritime, and the RNT Foundation.
- GPS Backup/Complementary PNT Demonstration
 - DOT conducted a demonstration of 11 PNT technologies last year, which was required by the FY18 National Defense Authorization Act. There was a wide variety of PNT technology vendor participation, including two low Earth orbit companies and a couple companies demonstrating time over fiber. They left it up to the vendors to choose which scenarios they

- wanted to demonstrate. They executed three field campaigns, technology demonstrations, and analysis and assessment of data.
- The work culminated in a report they submitted to Congress on January 15. (Along with NTRSA Report)
 - Observations from the FY18 NDAA PNT Demonstration
 - Results indicate there are suitable and mature private-sector PNT technologies that have the potential to meet a diversity of application-specific needs.
 - The demonstration was designed to showcase technologies in the “best light” possible, and complementary PNT technologies were not stress-tested. There is follow-on work to be done in terms of understanding vulnerabilities and limitations.
 - The transportation sector has some of the most stringent PNT performance requirements in terms of accuracy, integrity, and reliability. Not all safety-critical transportation requirements may be met by market-based business models for PNT technologies.
 - Private-sector complementary PNT technologies do not currently have the level of open specifications and standards that have made GPS such a critical and widely adopted service. A similar level of standards, resiliency and vulnerability testing, and performance monitoring must be developed for these technologies.
 - Complementary PNT: Recommended Next Steps
 - Safety-critical PNT requirements and standards development for transportation services
 - PNT vulnerability and performance testing framework for demonstrated and suitable complementary technologies.
 - Procedures, facilities, and platforms for testing PNT performance and resilience to threats
 - Certification protocols for safety-critical PNT functions
 - PNT performance monitoring capabilities to ensure operational PNT services provide resilience and achieve safety-critical standards for transportation and critical infrastructure applications.
 - DOT Focus on PNT for Highly Automated Systems
 - PNT for Automated Vehicles: Intelligent Transportation Systems (ITS) Joint Program Office
 - AV use cases / scenarios
 - Determine PNT requirements for AV operations
 - Assess GNSS and other candidate sensor technologies
 - Analyze PNT performance of individual sensors
 - Determine navigation performance enhancements achieved by sensor fusion
 - DOT University Transportation Center: Opened last year.
 - Vulnerability cataloging and test threat vector development
 - Resiliency testing
 - Standards, Guidelines, and Best-practices for cyber resiliency
 - OST-R Highly Automated System Safety Center of Excellence
 - Resilient PNT Services for Highly Automated Safety Systems
 - Focus on NIST PNT Profile
 - They have had a lot of spectrum challenges, particularly with GPS, but that is expected anywhere navigation capabilities are subject to interference. They are also looking at toughening user equipment, authenticating signals, and new antenna designs and algorithms.

Mr. Platt provided an overview of the National Risk Management and Program Strategy and how DHS uses risk management as a construct for PNT.

- Mitigation via Vulnerability and Impact Assessment: (Vulnerability Assessment)

- They have been focused on testing, beginning with characterizing receivers. They have a classified program that can look at GPS receivers, determine how vulnerable they are to jamming and spoofing, and determine how they behave when jammed or spoofed. Some handle it well and others simply accept the data that comes in. They are working with receiver manufacturers to build more secure and resilient receivers.
- Address vulnerability in systems: If improper data comes into a system, there should be other environmental data available to validate the GPS information and determine if certain processes should be allowed. One example is an aircraft that was using GPS to maintain level flight. There was jamming, and suddenly the aircraft was not maintaining level flight. That should have never happened because there are many systems in the aircraft to tell you whether or not you're on level flight.
- They have been working with all of the sectors to develop plans to conduct vulnerability testing to identify their most critical systems and determine if there are PNT dependences. How will those systems react if there is no PNT data available to them or if the PNT data is corrupted? GPS has been so reliable that it is always assumed to be right but that's not a good approach. We should treat a GPS signal as an untrusted source until it can be validated.
- Mitigation via Awareness (Engage and Educate)
 - Best practices
 - Need for resilient PNT equipment
 - End-users down-stream impacts
- Mitigation via Improved Equipment (Conformance Framework)
 - DHS has worked with inter-agency partners and the private sector to develop the Conformance Framework, which was released for public review. They are trying to categorize by level of security and resilience. At the bottom level, basically if something goes bad it can be reset. At the top level, it will see that there is either no data or anomalous data and it will continue to operate for an extended period of time without external input or with corrupted external input. Designing systems through the Conformance Framework will be helpful.
 - The long-term goal is to take it to IEEE or another standards organization. They don't want to tell manufacturers how to build to this standard. They will each develop their own methodology of complying with the framework. They will decide whether or not there should be an organization that can test the bodies and validate and if self-certification through the individual manufactures will work.
- Mitigation via Diversity (Complementary PNT)
 - They are looking at alternate PNT services, which include existing services that people have never considered. There are situations where people need a 1 second level of accuracy and they have access to a network, yet they still use a GPS receiver as the primary timing source for a particular application. Every additional PNT service you bring to your system brings a new opportunity to potentially get more resilience, more accuracy, and more robustness.
 - The goal over the long term is to bend the dependence on GPS and GNSS. Reasonable and prudent steps should be taken to seek out alternate sources of PNT and we should always be responsible in how we use them.

Discussion

- The Chair asked if they are they getting cooperation and a good response from vendors on the Conformance Framework?

Mr. Platt said yes. They realized it could potentially create pushback, so they intentionally brought the vendors, including many of the major manufacturers, into the working groups. They

have been telling owners and operators for years that there are vulnerabilities with the GPS systems. The owners and operators then start demanding better equipment because they recognize the risk.

The Chair thanked the speakers and recessed the meeting for a 14-minute break.

Board members received a presentation from the Department of Commerce Ethics Office on the “Top 10 Ethics Rules for Locally Engaged Staff of the U.S. Commercial Service.” Follow-up information and financial disclosure forms will be sent out. Mr. Scholl clarified that the rules are applicable when board members are serving in their capacity as SGEs.

Final Board Reviews, Recommendations and Discussions

The Chair invited board members to share input on the presentations and areas of recommendation.

- Ms. Hallawell commented that during the presentations she wasn't sure whether the questions she had were too narrow or too broad. Is there a standard set of questions for board members to ask?

The Chair said there are a wide range of issues that come before the board, and in some cases there is an obvious set of issues or concerns. The SolarWinds briefing is a good example. In other cases, the presentations are just informational.

Mr. Groman said that he thought about the same thing when he first joined the board. He has found that questions asked are rarely irrelevant or off topic. He encouraged engagement and exploration, which can provoke a good discussion. The federal government is very different from the private sector, so the diverse backgrounds and perspectives of the board members are important.

Ms. Hallawell said it might be useful if prior to the meeting they could ask what questions or feedback the speakers are looking for because the speakers put a lot of time and effort into their presentations.

Mr. Scholl said they can solicit that from the speakers when they're invited.

Ms. Hallawell commented that nearly every presentation involved testing of vulnerabilities. It could be useful to think about the major themes that come up in every presentation when offering guidance.

- Mr. Groman asked if there would be privacy events at NIST in 2021.

Mr. Scholl said he would check.

- The Chair said that Ms. Fitzgerald-McKay drafted a letter on secure software configurations. He asked about language in the draft addressing network configurations and software configurations.

Ms. Fitzgerald-McKay said she wrote the draft very broadly in the hope that the board members would help scope it better. The main point remains that vendors will know best the appropriate secure configurations.

Ms. Hallawell said it seemed to be focused on secure cloud network configurations, but instead of being too broad, is it too narrow?

Ms. Fitzgerald-McKay asked for suggested wording.

Ms. Hallawell asked whether the recommendation should address not just the configurations themselves but also the timing of availability to other organizations when there is some

intelligence of a breach or attack. Most people would put in place the configuration changes when they know there's an active attack. A lot of the issue is about the timing and the awareness of an attack.

Ms. Fitzgerald-McKay asked whether she is suggesting that they request not only a secure out-of-the-box configuration but also regular updates of the configuration in light of new attacks.

The Chair said it is better to go with broader language: a) Anything that can be configured ought to be shipped configured securely. b) If there are emerging attacks or threats, then it is a best practice for vendors to ship updates as needed.

Ms. Fitzgerald-McKay asked who the request goes to and who the requirement would apply to.

The Chair asked Mr. Scholl for history on previous configuration guidance. If it's technical guidance based on best practices, it is aimed at vendors. If it's procurement requirements, it's probably OMB.

Mr. Scholl said the board's scope is to provide advice to NIST, DHS, and OMB. For configurations, historically they've done all of it. It's not uncommon for an organization to have its own configuration policies. OMB has had configuration requirements. You can recommend that NIST work with partner agencies and industry to develop secure configurations. Procurement requirements would not necessarily be in the scope of the board. Requirements for use would be under OMB.

The Chair asked if the board is authorized to communicate recommendations to the director of NSA.

Mr. Scholl said he would double check, but he believes NSA is not in the board's scope.

The Chair said that the letter could go to NIST and DHS suggesting they work with partner agencies and industry to identify areas where the creation and promulgation of secure configurations would benefit the security of government systems and evangelize those configurations to the vendor community to supply and the federal user community to adopt. The board can wordsmith after the meeting, assuming the board agrees that they want to go forward with the recommendation.

Ms. Hallawell said they should offer dual-pronged advice so 1) federal agencies implement the most secure configurations and 2) vendors easily provide those configurations in a time-sensitive way. She suggested running a draft by DHS and pressure test it before sending out.

Mr. Scholl confirmed that the board needs to decide on the concept and approach and then can iron out the wordsmithing between meetings.

Mr. Gattoni said he could float a draft. He added that he really likes the Chair's suggestion that any bit of development product that ships has to be done securely.

Ms. Fitzgerald-McKay asked if there is a need to have some sort of machine-readable expression of this configuration being made available to the consumer to support things like testing and checking for updates.

Ms. Hallawell said there is a common thread of secure software in general that the board needs to make a recommendation about across different applications. Could they say ISPAB is working on broader guidance in this area and then get feedback? They want to be sure the guidance they offer is helpful.

Mr. Groman asked for clarification on who they are asking to do what. Are they suggesting the government work closer with the private sector? Have they identified a flaw in a policy, in a document?

Ms. Fitzgerald-McKay said there is a gap between available secure configurations and what ISPAB can do about it.

The Chair said that back in the 2008-10 era, there was a lot of emphasis by NIST and NSA on getting vendors to create secure configurations and secure configuration guidelines. In the intervening decade, the focus on secure configuration has diminished while the importance of having secure configurations – and having them by default - has increased. The board could try to get the relevant agencies, NIST, DHS, and NSA, to re-energize that work with industry.

Mr. Groman said he agreed with the principle and the concept, but he wasn't clear on exactly what levers they are trying to move. Is it a FedRAMP problem? Is it a procurement issue?

The Chair said there is a NIST and technical agency lever to get the momentum on secure configurations, and there may be procurement lever that follows.

Mr. Groman asked what incentive or disincentive is not occurring currently. What is the solution they are looking for? Where is it that they need a shift? Is it in a NIST document? Require something in the RMF?

Ms. Hallawell said the core problem outlined in the SolarWinds brief is that most organization do not implement the securest software configurations and do not always update on a timely basis. When a massive vulnerability or attack comes to light, everyone rushes to put the secure configuration in place. Are the vendors making the most secure configurations available by default? She suggested the problem is that they're not being used.

Mr. Duvvur said there are other angles to the issue. If you think about it from a software perspective versus the as-a-service model, you're increasingly seeing a shift-left approach to putting security at the front end of DevSecOps. There is a full lifecycle aspect to it. Are you ensuring observability of the configurations after they're deployed? There is a broader set of deployment and lifecycle management principles that need to be accounted for.

The Chair said it is a complicated problem. If the solutions aren't in the products, the agencies can never deploy them. If the solutions are in the products, then some agencies might deploy them. There is a difference between a vendor including it somewhere and a vendor including it by default. That is a big difference for the user. In the real world, that is better than a procurement requirement.

Ms. Fitzgerald-McKay asked if they are asking NIST, DHS, et.al. to investigate methods of incentivizing vendors to ship secure-by-default configurations.

Mr. Groman said he agreed completely with the principle, but he wanted to find a way they could truly influence the issue. Where is the pressure point? The issue with FedRAMP is the constant tension of how long they take to vet security versus the pressure to deploy.

Ms. Hallawell said it not necessarily just on the vendors. There is a mess across the board. Users don't always want to use technology in the most secure mode because it impedes usability. There's also drift. There are solutions that monitor how you're implementing security solution. The board can't be too high-level or too prescriptive in the recommendation, but they do want to raise the bar.

The Chair said he likes the idea of asking NIST and DHS to provide the board with feedback on ways to apply leverage to move to secure configuration across products that are configurable.

Mr. Scholl said they would be happy to provide that. That could be the initial recommendation directed to NIST and DHS: Come back and provide the board areas where the most assistance is needed in ensuring federal agencies are implementing secure configurations.

Ms. Fitzgerald-McKay said maybe they could drill down a little bit further to Mr. Groman's point regarding where in different policies there could be recommendations.

The Chair asked for a motion on producing a letter to send to the relevant addressees at NIST and DHS.

Ms. Hallawell moved that the letter be produced.

Mr. Groman seconded the motion.

The Chair asked for any objections.

Mr. Groman added that he has no objection, but the letter needs to show that the board has an understanding of the nuances and complexities of the issue. There is an enormous amount of flexibility left to agencies at present.

The Chair said final wordsmithing will be done after the meeting.

Board members voted unanimously in favor of the motion.

- The Chair said that two other letters were drafted based on the previous day's discussions, and they are subject to final revisions via email following the meeting.

Ms. Hallawell asked if the letter about security training should be very specific and targeted. What is the point of it apart from recommending making a new effort?

The Chair said the letter is to move NICE beyond focusing on the security workforce to the workforce that actually makes a difference in whether systems are secure or not.

Mr. Groman asked if it meant all people using networks.

The Chair said it is focused on people designing and developing systems and networks. The idea is to equip those building the systems and networks to build them securely.

Mr. Groman said it seemed like a straight-forward recommendation.

The Chair asked if everyone was comfortable with the two draft letters.

No objections were raised.

The Chair said he would revise the two drafts, and Ms. Fitzgerald-McKay would revise the letter she drafted.

The Chair thanked Mr. Scholl.

Mr. Scholl said the next ISPAB meeting is scheduled for June 23-24, 2021. A decision will be made no later than early May on whether it will be virtual or in-person.

ATTENDANCE		
Board Members		
Steve	Lipner	SAFECode, Chair, ISPAB
Brett	Baker	U.S. Nuclear Regulatory Commission
Douglas	Maughan	NSF
Akilesh	Duvvur	IBM
Jessica	Fitzgerald-McKay	NSA
Brian	Gattoni	DHS
Marc	Groman	Privacy Consulting
Arabella	Hallawell	NETSCOUT Systems
Phil	Venables	Google Cloud
NIST		
James	St. Pierre	NIST
Matthew	Scholl	NIST
Kevin	Stine	NIST
Elham	Tabassi	NIST
Katerina	Megas	NIST
Kim	Schaffer	NIST
Patrick	Grother	NIST
Craig	Watson	NIST
Jeff	Brewer	NIST
Caron	Carlson	Exeter/Hill
Warren	Salisbury	Exeter/Hill
Speakers		
Jay	Gazlay	DHS
Michelle	Mazurek	University of Maryland
Josiah	Dykstra	NSA
Frank	Nagle	Harvard Business School
David	Wheeler	Linux Foundation
Stacy	Bostjanick	DHS
James	Platt	DHS
Karen	Van Dyke	DOT
Registered Attendees		
Olatokunboh	Abereoje	FAA (ATO Cybersecurity Group)
Jill	Abitbol	Cybersecurity Law Report - Acuris
Eduard	Alpin	Verisk Analytics
Mariam	Baksh	Nextgov
Avonne	Bell	CTIA
Richard	Beutel	Cyrrus
Peter	Bloniarz	New York State Cyber Security Advisory Board
Donald	Claus	Lawrence Livermore National Laboratory
Dylan	Cohen	US House of Representatives
Steve	Conley	Wiley Rein LLP
Amanda	Coolidge	self
Olaf	Corning	The U.S. House of Representatives
Kathryn	Daily	BAI Information Security
David	Danks	
Jacoby	Davis	Coalfire Systems

Michael	Diakiwski	Wiley Rein LLP
Harry	Doyle	HD Healthcare, LLC
Patrick	Eddington	Cato Institute
Christopher	Egan	IBM
Rachel	Emmons	US House of Representatives
Charles	Faulkner	ICF
Andrew	Fausett	U.S. Senate Committee on the Judiciary
Matt	Fleischer-Black	Cybersecurity Law Report
Joelle	Foucher	BMC
Meeghan	Francisco	Cyber4Heroes Program- student
Kevin	Fu	FDA CDRH
Eric	Geller	Politico
Chad	Grant	GCG
Douglas	Gray	U.S. Farm Credit Administration
Eric	Grosse	
Kyle	Gutierrez	Wiley Rein LLP
Chloe	Hawker	Wiley Rein LLP
Jory	Heckman	Federal News Network
Manuel	Hernandez	Procapitol
Mat	Heyman	Impresa Mgt Solutions, LLC
Joseph	Hoellerer	Security Industry Association
John	Holmes	American Honda Motor Co.
David	Holtzman	HIT Privacy LLC
Daniel	Hucko	St. JOHN FISHER college
Katie	Ignaszewski	IBM
Jason	Jackson	GAO
Vince	Jesaitis	Arm, Inc.
Michael	Kans	Michael Kans Law
Jason	Kerben	Department of State
Filipp	Khosh	NUARI
David	Larsen	Cyber4Heroes Program- student
Tom	Leithauser	Telecommunications Reports
Sofia	Lesmes	
Ioana	Lewis	CompTIA
Frank	Limardo	Albertsons Companies
Sean	Lyngaas	CyberScoop
Mahlet	Makonnen	Williams & Jensen
Timothy	Matthews	INSCOM
Jay	Meier	Sage Capital Advisors, LLC
Christina	Morgeneier	
Derick	Naef	ReFirm Labs, Inc.
Ashish	Nangpal	VERISK
Bill	Newhouse	NIST/NCCoE
Shauna	O'Leary	BCS365
Hugh	Paquette	US GAO
Celia	Paulsen	NIST MEP
Kristen	Pedersen	Norwich University
Brendan	Peter	IDEMIA
Boris	Polania	Honda
Alex	Richmond	NorthPoint Data Security
Greg	Richmond	NorthPoint Data Security
Renault	Ross	RNSC Technology
Raymond	Savarda	Sensus

Fred	Schneider	Cornell University
Gabrielle	Shea	NEC
Barry	Skidmore	CISA
Annie	Sokol	NIST
Roberta	Stempfley	
Alisha	Stevens	DoD
Travis	Stoller	Wiley Rein LLP
Dwayne	Tanner	True North Cyber Solutions, LLC
Raj	Telwala	Touro College Jacob D. Fuchsberg Law Center Student Extern
Sandra	Throneberry	Lockheed Martin
Charles	Tupitza	Americas SBDC
Peter	Weinberger	Google Inc
Troy	Wells	CISA
Jake	Wiener	Electronic Privacy Information Center (EPIC)
Matthew	Wo	WE Comms
Mike	Wolbrink	Azule
Steven	Wray	CornerStone Bank NA