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

October 25-26, 2023
JW Marriott Washington DC,
The Senate Room (Lobby Level)
1331 Pennsylvania Ave, Washington, DC 20004

Board Members
Steven Lipner, SAFECode, Chair, ISPAB
Dr. Brett Baker, NARA
Giulia Fanti, Carnegie Mellon University
Jessica Fitzgerald-McKay, NSA
Alex Gantman, Qualcomm
Brian Gattoni, Federal Reserve Board
Cristin Flynn Goodwin, Advancing Cyber
Marc Groman, Groman Consulting
Arabella Hallawell, WhiteSource (absent)
Essye Miller, Executive Business Management (EBM)
Katie Moussouris, Luta Security
Phil Venables, Google Cloud

Board Secretariat and NIST Staff
Matthew Scholl, NIST
Jeff Brewer, NIST
Jim St. Pierre, NIST
Diana Proud-Madruga, Electrosoft//Exeter Government Services LLC

Wednesday, October 25, 2023

Welcome and Opening Remarks

Steve Lipner, Chair, ISPAB, Executive Director, SAFECode

- The Chair opened the meeting at 10:00 a.m. ET and welcomed everyone to the meeting.
  - Invited each board member to briefly share a few words.

Board Member Introductions and Updates

Ms. Miller

- The US Cyber Challenge (USCC) recently announced winners for 2023.
- The USCC is looking for participants for in next year's challenge.

Ms. Fitzgerald-McKay

- No updates from her team were provided but announced that the NSA will be opening an AI security center adjacent the current facility.

Ms. Fanti

- Currently studying how synthetic data can be used for various use cases, including releasing population-level analytics.

Mr. Venables

- Doing a lot of AI work, both security and risk management of AI and AI applied to cybersecurity.
- Google launched the secure AI framework which is trying to lift the conversation around how to do end-to-end risk management, trust, and safety of AI as opposed to just focusing on the micro cybersecurity challenges.
- On post-quantum cryptography, they are making more progress in deploying some of the pre-standard algorithms on their backbone infrastructure.
    - They're running Kyber as part of some of the key exchanges, which is going well so far.
- Currently a member on the White House President's Council of Advisors on Science and Technology.
    - Probably about six to eight weeks away from publishing a cyber physical resilience report and recommendations.
    - Will include a whole bunch of things around recommendations for sector specific risk management and various other things. Have also had various sessions with NIST.

Mr. Gantman

- Spending time thinking about measuring security outcomes.
    - Not coming up with anything
- Also spending a bunch of time reading various security literature.
    - Disappointed that there's not much actionable guidance that we can offer people on secure design.

Mr. Gattoni

- In a new job, updating a series of strategic documents.
    - It's the update year of the IT strategy and the cybersecurity strategy.
        - It's a good way to inject some of the stuff we hear here into the middle of that.
        - Gen AI is a big topic and figuring out where it can help and doing so in a secure manner is important.

Ms. Flynn-Goodwin

- No longer with Microsoft after being with them for 20 years.
- Founded two companies:
    - One that's focusing on incident response and threat intelligence.
    - Another that is a law firm that advises on threat intelligence and incident response.

Mr. Groman –

- Currently teaching threat assessment and incident response at Georgetown Law School.
    - Cybersecurity, for lawyers.
- Building on what Phil [Venables] said, he's constantly trying to inject into the discussion of secure AI and talk about safety, trust and responsible AI and try to broaden the concept of "secure".

Ms. Moussouris

- Her company is part of an industry association, the Hacking Policy Council, and she serves as one of the advisory board members.
    - Members include Intel, Google, Hacker One, Bug Crowd, and Integrity (the only European partner).
    - Sent a letter from the Hacking Policy Council about the Cyber Resilience Act (CRA), voicing some concerns about the vulnerability disclosure requirements.

Mr. Lipner

- Based on this, and other conversations he's had with some of the board members, it seems that ISPAB should be ramping up their focus on AI and ISPAB's position on AI.
    - Possibly make AI security a focus of the next meeting. Would like to spend some time brainstorming what would be interesting to talk about.
- SAFECode continues to be involved in the various software assurance initiatives and dialogues.
- Chair of a new national academies committee on assurance and resilience of large, complex government software systems, which will produce a future report.

ITL Update

Jim St. Pierre, Acting Director, Information Technology Laboratory (ITL), NIST

Purpose

- Stays the same as always: Cultivating trust in metrology.
    - This is the business focus and lens through which they look at their work.

NIST Wide Critical and Emerging Technologies

- These are defined by the Under Secretary for NIST, Laurie Locascio, as the NIST critical emerging technologies.
- ITL is involved with all of these through cybersecurity work and in our math and statistics division, but it's exciting that we're tied into all the NIST critical areas.

ITL Portfolio

- The portfolio is broken out into supporting those areas of critical emerging technologies through fundamental research.
- Computer science, math and statistics supply a lot of the basis for standards and guidance development.
- One of the challenges is to keep the right balance between applied work and fundamental research.
    - It's the pre-standards research that's needed to effectively participate in national and international standards.

Artificial Intelligence

NIST is interested in hearing the board's thoughts on the security of AI and AI for cybersecurity.

- NIST's AI program is focused on fundamental and applied research to lead toward standards.
    - Doing a lot of work on guidance, tools, and frameworks.
- AI Risk Management Framework
    - Informed by the Cybersecurity Framework (CSF) and the approach that was taken to the CSF including some of the core functions, such as Governance. This is a key element of the AI framework as well.
- NIST's history is as a measurement lab.
    - One of the greatest challenges and focus is testing, evaluation, and validation of AI models.
- NIST has taken a leadership role in AI.
    - The AI team has done an amazing job, led by Elham Tabassi, both domestically and internationally.

Generative AI Public Working Group

- Secretary Raimondo announced the AI Public Working Group in July, and it is off to a fast start.

  - Over 1000 volunteers have joined the group to help us work on a profile.
    - Similar to CSF profiles.
    - Same concept used with the AI Risk Management Framework.
- Looking at four areas:
  - Governance,
  - Digital Content Provenance,
    - Critical for things like elections
  - Pre-deployment verification and validation,
    - How you think about making sure you're testing and looking at your AI before you deploy it, not as you're doing it
  - Incident Disclosure.
- Generally the four are running their own schedules for anticipated release for public comment later this year on those four sets of guidelines.
- NIST can't do what we do without the volunteers that we get internationally.
  - Allows NIST to have a very significant impact.
- The Administration and Congress is interested in how to reap societal, commerce, and business benefits, while managing the risks.
  - That's why the AI risk management framework is such an important and useful tool in that effort.

National AI Advisory Committee

- NIST provides the Secretariat for the National AI Advisory Committee.
- There are nine areas of focus areas that were realigned recently:
  - AI Futures: Sustaining Innovation in Next Gen AI,
  - AI in Work and the Workforce,
  - AI Regulation and Executive Action,
  - Engagement, Education and Inclusion,
  - Generative and NextGen AI: Safety and Assurance,
  - Rights-Respecting AI,
  - AI and the Economy,
  - Procurement of AI Systems,
  - International Arena; Collaboration on AI Policy and AI-Enabled Solutions.
- The AI Advisory Committee has been very busy this summer:
  - Had six workshops this summer.
  - They're in their second year.
  - They've had a recent report on the recommendations for international and emerging economies and how to truly engage with them as the AI efforts move forward.
  - Three constructive explainer documents,
    - Overview of the AI RMF,
    - Generative AI, and
    - AI regulation.

NIST Leadership in AI

- Elham Tabassi is on the cover of Time Magazine as one of the 100 most influential people in AI.
- As part of the Singapore critical and emerging technology (CET) dialogue, a crosswalk between NIST's AI Risk Management Framework and Singapore's Model Governance Framework was published.

- NIST leadership hosted a workshop with ANSI on risk management of AI systems, and
- Held an AI Governance Forum in Seattle, Washington that was hosted by Senator Cantwell and Under Secretary Locascio.

What's Next for NIST

- A joint webinar on AI risk and safety with the UK coming up at the end of the month is planned,
- Continued engagement including a lot of international engagements.
    - Elham and her team are doing a lot of traveling.
- Measurement is critical for NIST so they're going to expand on their testing and evaluation standards and guidance.
    - We think those are core elements of AI going forward, including joint R&D with a stakeholder community.
- Looking forward to continuing to lead domestically.
    - There's an executive action expected and anticipate that NIST will have a significant role.

NIST – ITL and Biometrics

- ITL conducts biometrics work that supports law enforcement. Listed below are different evaluations of biometric measurements, where NIST provides evaluation expertise:
    - Nail to Nail (N2N) Fingerprint Capture Challenge,
    - Fingerprint Vendor Technology Evaluation (FpVTE),
    - Slap Fingerprint Segmentation Evaluations (SlapSeg),
    - Proprietary Fingerprint Template Evaluations (PFT),
    - Minutiae Interoperability Exchange (MINEX),
    - Evaluation of Latent Fingerprint Technologies (ELFT),
    - Biometric Quality.
- Working on testing the accuracy of these different approaches.

NIST – ITL and Face Recognition

- NIST's facial recognition work has received a fair amount of attention.
- A recent report on bias and facial recognition has generated much interest from Congress.

NIST – ITL and Information Retrieval

- Three workshops planned for November, 2023
    - TREC workshop at NCCOE - https://www.nist.gov/news-events/events/2023/11/trectrecvidtac-2023/trec
    - TRECVID workshop - https://www.nist.gov/news-events/events/2023/11/trectrecvidtac-2023/trecvid
    - TAC workshop - https://www.nist.gov/news-events/events/2023/11/trectrecvidtac-2023/tac

NIST – Quantum Information in ITL

- The Joint Center for Quantum Information and Computer Science (QuICS) is a partnership between NIST and the University of Maryland that advances research and education in quantum computer science and information theory.
- NIST also partners with one of the nation's leading research institutes in the physical sciences. The Joint Institute for Laboratory Astrophysics (JILA) conducts research that addresses fundamental scientific questions about the limits of quantum measurements and technologies:

  - Focused on the physics of quantum computers. How do we build the physics?
  - Focused on the information and computer science aspects.
- NIST Statistics:
  - 14 fellows: 8 NIST fellows, 20 postdocs and 68 students.
- Publications in 2023:
  - A recent paper on quantum algorithms and the power of forgetting, finding that the quantum algorithms could find their way through a maze better if they didn't keep track of where they've been.
- NIST is taking on leadership roles in DC-QNet, quantum network.
  - A regional quantum network.
    - Joint effort between NIST, NASA, NSA, Army Research Lab, and Naval Observatory
    - Able to send single photons.
    - Operational for research on things such as precise timing for synchronization and noise characterization.

Encryption Updates:

- Three New Federal Information Processing Standards (FIPS) Drafts have been released.
- NIST thanks Phil Venables for backend testing on Kyber.
- Names of the algorithms will need to be changed for the standards, but current names are expected to remain when referring to them conversationally.
- Intention was to do Kyber and Dilithium first. Sphinx+ is in a finalized state, it didn't have a lot of different parameter options that we were debating about and so it was ready to include. Will publish Falcon next.
- NIST is working on an open call for onramp signatures.
- NIST SP 800-208: Describes a stateful hash-based signature method that's primarily used in case someone needed to immediately implement a code signing capacity.
  - It uses current hashes, so it didn't need a FIPS, but it is brittle, and it is not misuse resistant.
  - Haven't been publicizing it outside of very niche use cases,
  - Based on feedback from industry, NIST may be updating or looking at other potential methods to use the stateful hash based signature that may provide better misuse resistance while allowing for backup at the same time,
  - NIST may be open for discussion or have a workshop on that in the future.
  -
- **Mr. Venables –** Asked how the partnership with other countries is going, particularly European countries, on whether they're going to adopt these standards?
  - **Mr. Scholl –** Replied that they are pleased with it. It's not necessarily uniform.
    - Many of the EU countries individually have said they are going to follow and use our standards.
    - UK is not part of the EU, but UK has also said that. There are some individual EU countries who have said they also intend to use other algorithms for their internal national algorithms.
    - Germany has been talking about FrodoKEM.
- **Mr. Venables –** Asked a similar question regarding Asian Pacific countries like Japan and South Korea?
  - **Mr. Scholl –** Japan, South Korea, Taiwan, and Australia are pretty much all on board.
    - Having a meeting with Taiwan next week.
    - ISO SC 27 meeting was two weeks ago. They brought the FIPS drafts to ISO SC 27.

National Cybersecurity Center of Excellence (NCCoE) Activities

- NIST and the state of Maryland renewed their partnership in support of the NCCoE. It's a five-year partnership intermediary agreement continuing the collaborative effort where industry, government, and academic experts work together to solve pressing cybersecurity challenges.

NIST – ITL and NCCoE Projects Underway

- 5G Cybersecurity
  - Creating practical solutions to improve the security of a system's architecture, provide a secure cloud-based infrastructure, and enable security features provided by 5G standards.
  - Have a 5G network back end simulated with equipment from participants.
- Cybersecurity Framework to secure the infrastructure for extremely fast charging vehicles.
- Cybersecurity of Genomic Data
  - As it becomes easier and cheaper to sequence genomes, all that data needs to be protected. ITL is working with the biologists at the Material Measurement Laboratory (MML), on developing some additional guidance on securing genomic data.
- **Mr. Groman –** Asked what would be done differently about cybersecurity for genomic data versus other data? As technology evolves, it becomes easier to identify genomic data and link to humans, increasing the value of the data and therefore the sensitivity and risk. But the cyber security principles are likely the same as with other important data.
  - **Mr. St. Pierre –** Replied that he thinks the principles are the same.
    - We'd have to get one of the folks working on that. He hasn't been briefed recently on the work. It's a unique use case and guidance for that community to make sure they're aware of when they need to be thinking about security.
  - **Mr. Scholl –** Added that his understanding is that some of the technologies and storage capacity used is different than general IT. The concepts are the same but there's also a different technology base that might need different application than general purpose.
- **Mr. Lipner –** Mentioned there's an NCCoE project focused on the implementation of both migration and Secure Software Development Framework (SSDF). How are those progressing?
  - **Mr. Scholl –**Replied that he's not working on the SSDF. He's working on the migration, which is running. There are also several others such as CMVP.
  - **Mr. Lipner –** Mentioned that he thinks it'd be good to hear about those when they get to the point where NIST thinks it's appropriate.

NIST National Vulnerability Database (NVD)

- New API's have been added to the NVD.
  - Should make it easier to use and integrate for folks and to get access to the data.
  - Also using software IDs and linking to the NSRL; working toward making it easier to use and integrate with other tools.
- **Mr. Scholl –** Mentioned that they've updated the CVSS, the common vulnerability scoring system to a new version.
  - Added some contextual capacity for people who want to do local scoring to add their own mission or context.
  - Includes safety as a potential impact parameter.

  - Focused on the technology impact potential, rather than data of an event; sometimes the information does get misunderstood or misused. It doesn't necessarily correlate to other types of incidents.
- **Mr. Gantman –** Commented that it doesn't necessarily convey risk.
- **Mr. Scholl –** Replied that, to some extent, it is one part of an overall risk equation, but it is not a risk. It is information in case of a vulnerability.
- **Ms. Moussouris –** Commented that what she's heard about the NVD, around practitioners, is that there's often a lag in updating, thinking that is a resource allocation problem that could be solved with putting additional resources on it.
  - She has seen some research comparing NVD to China's equivalent and how much faster they (China) are able to update their National Vulnerability Database, quickly but selectively.
  - The main concern here is how many tools used by practitioners rely on consumption of NVD data. That's a reason to add more resources, to be able to provide more timely updates.
  - **Mr. St. Pierre –** Commented that they're always interested in looking at performance. If there's a need there, we need to have that conversation.
- **Mr. Gantman –** Asked if there is any discussion about adding more data about exploitation to the NVD? Right now, it contains one bit of information that says the vulnerability was exploited and there's no way to identify if it was one device via one endpoint exploited or 100 million endpoints exploited.
  - **Ms. Moussouris –** Commented that it's her understanding that the KEV was originally designed to show what is known, exploited in the federal space, then it was made public.
    - Now it has become the signal of what to patch first, which might not be applicable to a particular situation. It can be misleading about prioritization.
  - **Mr. Gantman –** Added that from the vendor perspective, if you know something about the vulnerability in the product we're producing and how it's being exploited, the vendor would like to know that too. Just having a single bit is not enough information. It would be useful to know how much NIST is working with CISA on that.
  - **Ms. Moussouris –** Mentioned that it might be useful if we were to give feedback saying is this (exploitation) only from the federal space that you are observing this, because the federal space often will lag far behind in patching.
    - Some of the first ones that appeared on the KEV list were 10 years out of date for patches,
    - Everybody else should be more up to date than that.
    - Is it scoped to the federal space, is it US only, is it global? How are they sourcing the KEV data?
    - **Ms. Flynn Goodwin –** Asked how is CISA looking at where they're seeing vulnerabilities in the federal space? Do they have a way to validate and articulate that? Asked if ISPAB wants to take a "we recommend" position?
    - **Ms. Moussouris –** Added they might only be limiting it to technology that's deployed in the federal space so understanding better how they've decided to expand that sample list would be good.
    - **Mr. Lipner –** Added that he has questions on the use of software ID (SWIDs) tags. He remembers discussions about the difficulty of figuring out what products and versions were being talked about in the context of vulnerability response and would like to know if SWID tags have proven to be the acceptable mechanism now.
    - **Mr. Scholl –** Replied that they don't know if that's the answer at this point. Ms. Moussouris was talking about timeliness and resources.

        - As of this data collection, there were 110 million individual CPE names.
        - Most of their time is spent generating CPEs, clearing up CPE, or figuring out what a vendor is calling their things so that we can have a good CPE to identify. SWID was an attempt at another method to help get better understanding of what is affected.
        - They're still working on that.

NIST – ITL and Workforce – NICE

- The Cybersecurity Career Ambassador Program
    - Working to create a network of volunteers to serve as champions for expanding and diversifying the cybersecurity workforce. Have an open call for proposals for the 2024 Nice Conference and Expo. Theme is strengthening ecosystems and aligning stakeholders to bridge the cybersecurity workforce gap. Workforce gap is a significant issue.
- Notice of Funding Opportunity (NOFO) for the RAMPS, regional alliances and multi-stakeholder partnerships to stimulate Cybersecurity Education and Workforce Development closed on September 5, 2023.

Cybersecurity Business Development Mission – Sept. 18-26, 2023

- Mr. St. Pierre went on a Business Development mission with the Undersecretary Laurie Locascio to Taiwan, and then was with the Deputy Secretary, Don Graves in South Korea and Japan. Goals were:
    - To introduce US firms to East Asia. They had 15 IT companies with them. Some were companies that are involved in communication technologies and critical infrastructure, and members of the global Critical Infrastructure Protection Market.
    - To assist in finding business partners.
    - Promoting NIST work such as the Cybersecurity Framework and other guidance and standards globally.
        - All countries expressed a lot of interest in NIST work.
        - Taiwan plans to translate the CSF 2.0 into Mandarin. Will work with them to make it an official NIST, US government translation.

Statistical Engineering Division

- Footwear Impression Comparison System:
    - Initial release of NIST Footwear Impression Comparison System that is more capable than many professional footwear examiners, with support from the FBI and other law enforcement agencies.
    - Statistical engineers are working to help improve the accuracy and the statistical analysis to support courtroom use.
- Certification of glycans that are used for testing of monoclonal antibody therapies. Currently working with biologists at NIST for statistical analysis and uncertainty analysis.
- Artificial Intelligence
    - Statisticians also are developing and continue to expand their ability to assist other labs with artificial intelligence related work, such as the one on the far right of deep generative modeling for comms systems testing and data sharing.
- Our statisticians have a long history of working with other NIST laboratories.

Staff Recognition

- NIST has a connection to the Nobel Prize this year in physics for experimental methods that generate attosecond pulses of light for the study of electric dynamics in matter.

- Department of Congress gold medal is highest award within the Commerce Department.
    - Won by NIST's Phish Scale team that developed a way for rating the difficulty of phishing attacks so they can better assess how you're performing and how your staff are picking up on training,
    - Another gold medal that our team was involved in the development of cell characterization standards to improve manufacturing quality of various lifesaving therapies, with our biologists,
- A silver medal for analyzing the effects of alkali-silica reactions for reinforced concrete.
- ITL has statisticians involved in many of the efforts that go on in our engineering lab looking at facilities, such as the:
    - National Construction Safety team that is looking at a variety of disasters and trying to inform federal building codes.
    - Morphing work as part of the Face Analysis Technology Evaluation (FATE) MORPH received a provisional patent for a methodology for detecting morphs using one-to-many face recognition.

Questions

- **Mr. Venables –** Asked, regarding NIST's own workforce challenges, if they are having difficulties finding the ability to hire and retain the right people versus all the pressure to pull people out of scientific establishments into industry.
    - **Mr. St. Pierre –** Replied that's an ongoing challenge. They've been pursuing authorities to be able to pay more and looking at ways to do that. One of the things that helps is that the work we do is very exciting. It's still a challenge trying to get people to come from industry.
    - **Mr. Groman –** Added that there's a huge interest in work in the government and places like NIST and cynicism about some of the industry jobs they are leaving.
        - The HR side of the house is where there's a challenge. It becomes a 10 month process and it's completely opaque. NIST has lost students who want to work in government, who are highly talented and go to the private sector, not because they want the extra $. HR just never got back to them.
    - **Mr. St. Pierre –** Commented that 60 to 90 days is the shortest. We were just talking about this at the Physics Committee for Advanced Technologies. It's a continuing challenge, but it's larger than NIST or even the Department of Commerce, He heard some encouraging news that their head of HR thinks he has some ideas on how to significantly speed things up.
    - **Mr. Groman –** Asked if most NIST positions require a clearance?
    - **Mr. Scholl –** Replied that most of the positions in his division do not.
    - **Mr. Venables –** Added that he would be happy to amplify NIST's postings. Just send them to him.
    - **Mr. Scholl –** Thanked him and mentioned they also put them on LinkedIn.
- **Ms. Flynn Goodwin –** Regarding facial recognition work.
    - There's been a lot of attention paid to China's advancements and is curious how involved the Public Company Accounting Oversight Board (PCAOB) is in some of the oversight of NIST's work in law enforcement and facial recognition?
    - Assumes they are engaging thoughtfully, thinking carefully, but then it hits the media and civil society suddenly puts you on a back foot because it looks worse than it may be. Asked if they've done any defensive thinking about privacy reviews or civil society reviews?
    - **Mr. St. Pierre –** Replied that they try to. They take comments very seriously; in privacy or bias, for example.

        - In that kind of discussion it is important to point out that what we're doing is critical to understanding the technology so that good policy decisions and recommendations for how the technology is applied or should not be applied can be made.
    - **Mr. Scholl -** We have not worked directly with the PCAOB on that project that I'm aware of. I'm not sure how active the PCAOB is or is not at this point.
- **Ms. Flynn Goodwin –** Added that she's thinking about where they can get an opinion that if NIST isn't acting, it's increasing the load from a privacy perspective. That NIST is looking at the technology from a neutral perspective to enable policymakers to have that conversation with the public.
    - Maybe get some privacy evaluation or opinion from the Privacy Office viewpoint.
    - **Mr. Groman –** Mentioned that he hasn't heard civil society direct their criticisms or concerns towards NIST's work as much as implementations.
    - **Ms. Flynn Goodwin –** Commented that a legal review of the statement and holding statement is always good too. Read a lot on surveillance technologies, particularly in Europe and in China.
        - If NIST is moving into this space, this is where you're going to see law enforcement moving more into this.
- **Ms. Fanti –** Regarding the facial recognition experiments, one of the things we've seen with many fairness metrics is that to achieve it, you must take a hit in accuracy. Wondering if they have seen that and if NIST has some structured way of thinking about how you trade off downstream performance with satisfying fairness or equity,
    - **Mr. St. Pierre –** Replied that he'd be happy to get her the report. His recollection is that the accuracy correlated with less bias.
    - **Mr. Scholl -** Added that the program runs tests and provides test data, and then exposes capability and capacity, thinking that was seen but not for all accuracy cases.
    - **Ms. Fanti –** Commented that she would be interested in hearing about the barriers to adopting un-biased machine learning models.
        - In a lot of machine learning models, if you try to enforce some kind of fairness, it typically causes you to take a hit in terms of overall accuracy, would be curious to hear if NIST has seen that trend.

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

Cybersecurity Threats Facing the USG, A Briefing

Colin Soutar, Managing Director, Deloitte & Touche LLP

Introduction

- Presenting some of the threats seen from Deloitte's perspective.
    - Threats that are evolving through our awareness of the commercial accounts and how we think that will impact the federal clients that we have as well.
- Deloitte is fortunate because about half of our cyber practitioners sit on the commercial side and about half on the government and public services side, which is federal, state and local. This allows Deloitte to pull across the two domains when doing polls around where things are looking around IoT or quantum cyber readiness. Roughly 5,000 in total.
- Background:
    - Was the CTO of a public company before he joined Deloitte 10 years ago.
    - Specialized in biometrics and identity technologies 20 to 25 years ago.

  - After the tragic events of 9/11/2001, helped NIST establish biometric standards, both nationally and internationally.
- Want to understand emerging tech and how that is evolving within the axes of risk, regulation, and technology. Where does regulation fit into that?
- Agenda:
  - Will talk a little about operational technology, and IoT.
    - Operational Technology is something that is aligned in an enterprise and can be managed as if it was an IT device. It's more of an enterprise technical deployment.
    - IoT is more commercial grade, internet facing devices that can be managed or manipulated accordingly.
    - Zero trust.
    - Post-quantum Cryptography
      - He's been running that aspect for Deloitte for the last three or four years.
      - Will get into the evolution, the timing, and the vendors who've been integrating some of the post-quantum cryptography algorithms. Where are the standards going and how that affects downstream activities around frameworks and regulations.
    - Generative AI
      - That's a topic that's surfacing a lot.

Operational Technology (OT)/Internet of Things (IoT)

- **OT Security Vulnerabilities**
  - Will start by talking briefly about OT and IoT.
    - OT is sometimes called cyber physical systems. What is the relationship to the overall enterprise risk profile that an organization has and how can you mitigate the threats in a harmonized way?
    - IoT is working with individuals/practitioners that are deploying technologies within some of the critical infrastructure, but their background and the reasons for deployment may not be in line with some of the IT security considerations.
      - Often, we're trying to understand the impact of those devices; what would happen if a denial of service occurs on a device or if the device is leveraged to get access to other parts of the enterprise?

- **Security testing of IoT Products Provides Immediate Benefits**
  - Why care about security for devices that are just a convenience?
    - Still see that a lot with our commercial clients. Some of the material that is put out is helping with that.
    - Believes that informative guidance typically works better.
    - IoT vendors could do with more prescriptive guidance because, otherwise, there's a lot of convenience aspects that they might want to evolve that are not necessarily in line with some of the security aspects that we hold true.
  - **Mr. Groman –** Asked if he could give a more concrete example of where a vendor may not be optimizing for security and where there's a need for more prescriptive guidance?
    - **Mr. Soutar –** Replied the one that came to mind is where they had hosted an event and were talking with a vendor about a traffic light system where they were gathering all sorts of data. They were asked what's the purpose for this data? The answer was "it might be useful in the future." It's the unbounded collection of data and then the aggregation of that which could then present a security issue.

    - **Mr. Groman –** Responded that makes sense but being more prescriptive is difficult, other than saying, "if you don't have an identifiable use for data, don't collect it" or providing updated minimization protocols and procedures. You can't remove the requirement for people to think.

Zero Trust (ZT)

- Everyone is asking, "Can I buy some of that zero trust? It's a concept based on principles. It's not something you buy off the shelf.
- Biggest challenges:
  - To do it properly, it can't just be thought of as "I'm going to substitute one thing for another." You need to change the overall governance. It needs to be part of the transformation from a cost effectiveness perspective, and you need to be able to do it in a way that is holistic.
  - Both federal and commercial clients come to us looking for steps they should take rather than working through a methodical roadmap that explains this big shift.
- **Ms. Miller –** Asked if they are finding that's driving more discussions on the underlying need for identity and credential management?
  - **Mr. Soutar –** Replied it is. The need to identify individuals and items, as well as objects, is driving that and getting more attention on it.
    - We may not be seeing the direct connection because it includes governance, strategy, identity, infrastructure, and so there are different ways of speaking about that.
    - Need to talk about what is the use case that we're trying to solve here? What is the outcome that we're trying to achieve? How do you get to that outcome? You need to have stronger authentication, better connectivity, better authorization, criteria that are policy, etc.
  - **Ms. Miller –** Mentioned that we've not done well with identity and credential management, and she was wondering if zero trust is driving a chance for that discussion, to force us to look at it differently.
  - **Mr. Soutar –** Replied that he thinks it is but does not have a concrete example. The dialogue is changing around it and the need for stronger credentials. FIPS 201 has had good success in dissemination and adoption.

Post Quantum Cryptography

- Threat to organizations may already be here. Problematic areas are:
  - People conflate timelines.
    - People think that they don't need to act until there's a more than a reasonable chance that a quantum computer will be here in the next two or three years even when presented with information on doing something like implementing Shor's algorithm, and the effect that it would have when implemented on current day public key cryptography.
    - The US government has done exceptionally well to get ahead of that, charging this back in what 2007 or 2006 to start with the algorithms.
  - Harvest now, decrypt later:
  - We don't know when this is going to occur.
    - There is a likelihood that a cryptographically relevant quantum computer will exist in the next 10 years, but we don't know the timeline.
    - The government has articulated extremely well in the National Security manuals that if you don't take those steps to do the inventory and look at what your exposure can be, then you

    - don't know how long it's going to take to upgrade all your infrastructure, your third-party dependencies, and understand exactly what you need to do.
    - But the impact is going to be massive across a multitude of different agencies and commercial suppliers. Some vendors are getting ahead of it.
  - What is going to drive the fundamental change?
    - On the federal marketplace are the federal agencies through directives.
    - On the commercial side:
      - We're concerned that it might be too late by the time they look at the standards and start implementing them.
      - We don't know what that pathway might be or not be to potential downstream frameworks that use the regulation or that's going to motivate some commercial practitioners or organizations to start to adopt post Quantum.
- **Mr. Venables –** Mentioned that when the NIST standards are published, regulators around the world in critical infrastructure sectors are going to stipulate an implementation timeline. The question is whether that's 1, 2, 3, or 5 years. So, it almost doesn't matter what day you believe a cryptographically relevant quantum computer exists, the actual deadline is going to be the time that the regulators are going to demand.
  - The US government is demanding implementation timelines from major technology providers.
  - The broader question is how much are you finding your customers realizing that the breadth of this change is not just a "rip and replace" a layer of crypto? It's going into applications and application protocols, for which keys and signature sizes are going to have to change.
  - They will have to change a whole bunch of stuff that isn't just cryptography. Are people realizing that from your experience?
  - **Mr. Soutar –** Replied, that he thinks there's not a deep enough or broad enough understanding, especially as it relates to the critical infrastructure.
    - The financial services industry is active in this realm now. If there were regulators that would pick it up earlier than others, it's likely to be in that sector.
    - The way we see the awareness maturing is from federal to the financial sector and then moving into other sectors like healthcare.
    - Is it our problem or are the Google's and others in the world going to deal with that for us? That's a little concerning regarding people embracing the problem.
    - Interested in ISPAB's perspective because he believes that, when FIPS standards are out, regulators may not pick them up as normative references.
  - **Mr. Groman –** Replied that maybe they're going to wait for it to go through ISO, or to be picked up in a framework like the CSF?
  - **Mr. Venables –** Added that he thinks it will be easy for them to put in their annual updates, "all regulated entities must submit a plan by X that shows by X plus three years" that they've adopted these things. He thinks we, as an industry and a society, have massively underestimated the high level application work that's going to have to happen.
      - For example, to find all the places where things like a signature block has now gone from 1k to 10k, or that requires them to change applications and every system like that. The thing that's going to take years is the crypto swap-out.
  - **Mr. Soutar** – Commented that they've been helping to try and drive messaging through the World Economic Forum for the last couple of years. He felt it was important that the right degree of messaging was out there. Don't want organizations federally and commercially to wait.
- **Ms. Flynn Goodwin –** Mentioned that interesting work in the threat and risk space is going to get created by the gap between those who implement and those who implement in part.

  - This will become a hygiene issue. What does that exploitation space look like? Have you seen anybody doing any interesting studies on when we make this migration, what that exploitation space looks like and then how does that create this whole new area of attack surface?
  - **Mr. Soutar** – Replied No, he has not seen such work. That's part of the concern, there's not enough understanding of even partial implementation adding exposed areas. There needs to be more discussion and a lot more analysis to determine the impact.
  - **Mr. Gantman –** Added that there's also the risk angle of just switching to freshly existing crypto implementations. It's unavoidable that we're probably going to introduce quite a lot of risk on our way to addressing the risk of quantum computers.
- **Mr. Soutar** – Commented that there are things like crypto agility and running in hybrid modes so that there's not an immediate flip and there's not a dependency on the new algorithm while it's being tested along with other things.
  - **Mr. Gantman –** Added, like all the software libraries.
  - **Mr. Lipner –** Commented that people are looking at being conservative and keeping both the old and implementing the new. You want to make sure you put those things in series, not in parallel. There's a variety of ways to screw up.
  - **Mr. Gantman –** Added that putting it in series helps when we look at the cryptographic text, but it doesn't help with like memory corruption attacks right?
  - **Mr. Lipner -** Added that a lot of the risk just comes from implementation complexity and implementation problems. People try to make massive changes in a big panic. That's not the usual prescription for quality.
  - **Ms. Fitzgerald-McKay –** Commented that for the national security systems problem, something we're concerned about is any sort of hybrid implementation that forces us to use both types of cryptography. It needs to be done in a way that, as folks are more confident, they can simply transition without being locked in for decades.
  - **Ms. Moussouris –** Added that an important part to flag for anyone who's going to attempt doing a hybrid or parallel support is that's a recipe for downgrade attacks. Should call that out explicitly and warn that this exposes your users to downgrade attacks and is not recommended.
  - **Ms. Fitzgerald-McKay –** Mentioned that there's a lot of geopolitical tension with the whole concept of quantum in general. The hybrid conversations are valuable but have the potential to open room for concern from folks who want us to be concerned about doing anything to prevent quantum cryptographic attacks.
- **Mr. Gattoni –** Commented, regarding the harvest now, decrypt later problem, system owners hear the threat part of it and come to the table asking how to protect themselves. In some cases they may not need to because it's data that will be publicly released before it can be decrypted anyway. This threat signal is clouding their investment judgment.
- **Mr. Soutar** – Added the question that if data were stolen two years ago and you know that they were stolen, then at the end of 2024, FIPS are released and there's an obligation through a regulator to use that, is that a go-forward requirement? Are you then responsible for stuff that's been stolen?
  - We're hopeful that there might be two byproducts:
    - There's lots of good federal guidance around protecting data and only acquiring what you need to acquire. On the commercial side, maybe it'll make people think twice about collecting data that they don't need because it's going to have a risk associated with it directly.
    - Quantum is a lens through which we've gone in and done assessments of people's cryptographic management capabilities.

- **Mr. Venables –** Regarding store and decrypt later: asked if any of their customers are starting to process the risks of storing and repudiate later. Particularly in areas where that's a high-risk issue?
- **Mr. Soutar** – Replied that he hasn't seen a lot of that yet.
  - There's been a lot of excitement and interest around Gen AI as a method to help what we're doing from a security perspective. Then, the other side of it is, the attacks that it may help, phishing and vishing, models that are relying upon the data provenance.
  - Starting to see lawsuits around IP, music, and provenance. I think that over the next few years, that's likely to be something that will only get greater as the reliance on data for making decisions becomes more important than custody. The ability to resist a quantum-based attack becomes part of that based on expectations you have for that data to remain secure, confidential, but also to have integrity.

Generative AI

- Gen AI has not been around long and is bringing threats and benefits as well. The types of threats (infiltration, expansion, and exploitation & exfiltration) that we're seeing. One of the biggest threats as we move forward is regarding data provenance. It seems to be overlooked a little bit.

Discussions

- **Mr. Scholl -** One point with the PQC. We have a good date when we think the new FIPS will be in. We haven't started our discussion about when to pull the RSA out, which is the other reflective side.
- **Ms. Flynn Goodwin –** Asked are these areas where you've seen attacks already, or are these hypotheses? I've heard limited use of AI attacks. We've seen a couple of reports. But I'm curious if you've seen all of these or some?
  - **Mr. Soutar** – Replied that he would poll the team and get back to them. Suspects it's probably 50-50 with some of them being hypotheses. Trying to develop an understanding what those threats are, not only nationally, but globally, too.

The Chair recessed the meeting for a 60-minute lunch break.

IoT Cybersecurity Guidance, Recommendations and Labels National Cybersecurity Implementation Plan

Kat Megas, IoT Cybersecurity Project Lead, NIST

Executive Order 14028

- The last time she spoke to the board was right after the executive order had come out. Were still in the process of just kicking off the labeling effort under the executive order.
- What we were asked to do under the executive order:
  - Develop and identify criteria that could underpin an IoT cybersecurity labeling program.
  - Work with other agencies to look at existing programs in the federal government and see what we can learn from those.
  - Pilot a labeling effort around both software labeling and IoT product labeling.

Elements of the Task

- Looked at the work in three parts:
  - One was identifying what the criteria are, what does the product or the product manufacturer have to do to support IoT cybersecurity?
  - The second part of it was how do you demonstrate that you are meeting those criteria?

  - The third one being how do you then communicate to audiences, what you are doing and what you've done to support IoT cybersecurity,
- The intent of the labeling scheme was to provide transparency to the customer so that the customer can make informed decisions when they're purchasing IoT devices.

Existing Non-sector specific IoT Security Baselines and Guidance

- Background
  - Started in 2017.
  - One of the initial areas looked at was trying to understand how we could help organizations that were adopting IoT manage their privacy and cybersecurity risks.
  - Published NISTIR 8228, looking at both the cybersecurity framework and the NIST RMF (the privacy framework had not yet been published).
  - Tried to articulate for organizations, how introducing IoT into your organization introduces new risks, how it may provide challenges, and how to address both privacy and cybersecurity risks.
  - One of the things we identified during those workshops was a disconnect between IoT manufacturers and users.
    - We asked organizations that were using IoT how they are mitigating some of the risks. Response was that they don't really have a good way to mitigate the risks. In some cases, they can't implement desired cybersecurity controls because it is dependent on some of the functionality that the product provides.
  - Started talking about how we can provide guidance for the manufacturers of products, helping them understand the minimum cybersecurity capabilities that they need to provide in the product.
    - Selected the word capability. Spent quite a bit of time talking with our partners and the private sector talking about the word because we felt that cybersecurity is a shared responsibility. The product provides the capability that the organization that's adopting the product makes use of.
  - Was around the time Executive Order (EO) 13800 came out. EO 13800 was looking at the botnet threat and when Commerce and DHS worked on the botnet report.
    - They identified IoT devices as one of the main threat vectors for botnets.
    - NIST was asked to develop a core baseline identifying the minimum cybersecurity requirements.
      - The reason we call it core is because it's not specific to any market,
      - Were asked to look at IoT devices.
      - Community of interest said that it was very important that we don't jump into this and cause fragmentation.
      - The core baseline has provided the foundation for both the federal agency and the federal government baseline as well as on the consumer side.

2019 Core Recommendations

- The core baseline was first published in 2019. Tried to balance the concept of providing a minimum and the interest of understanding that we don't necessarily want everybody to just aspire to the minimum,
  - Combined the baseline with a broader framework where we step the manufacturer through designing the product with the customer and the customer needs in mind.
    - When planning for cybersecurity, you have the baseline on one hand and you understand your customer, your customer risks, and what type of capabilities your customer has. Are they a

mature organization? What do you think they are going to be expecting you to provide in terms of cybersecurity in that product.

- The baseline was a component of this overarching process.

Core Baseline

- The core baseline consisted of two parts: technical and non-technical.
    - The technical component-
        - Developed by looking internationally and domestically at where we saw a lot of consensus from our partners and international standards.
            - What is minimum cybersecurity that any product should be able to support?
            - Are you able to identify the asset?
            - Do you secure the data? Do you secure the data in transit?
        - These are all very typical types of features that you would expect out of any software product.
- Non-technical capabilities-
    - More about the organizational processes in place that support the cybersecurity of the product throughout its lifecycle before you sell it to the customer and after the customer has taken control of that product. Covers things like:
        - Documentation
            - Software development plan?
            - Vulnerability management plan?
            - Software supply chain risk management plan?
    - Hoping to get manufacturers to start thinking about all the processes they need in place, not just technical capabilities.
        - Ways to hear from the community and gather information from the community about the cybersecurity of your product?

NIST SP 800-213 and 213A

- Anticipated that manufacturers would look to standards and best practices.
- The first time we adapted the baseline was coming out of the IoT cybersecurity Improvement Act of 2020. Were asked to build on the existing baseline but tailor it for federal government agencies so that when federal government agencies procure IoT devices, they had a baseline.
- NIST published 800-213 and 800-213A,
    - Federal agencies, per an OMB memo (December 2022), should be following the guidance in 800-213 and 800-213A, if they are procuring any sort of IoT device, as part of a product.
    - SP 800-213 provides the basis for federal agencies to think about the risk being introduced, the other components of the product, and where are the boundaries of the system.
    - Also gives them a catalog showing the minimum and, based on risk assessment, additional requirements beyond the baseline.

Consumer Baseline

- Held multiple workshops and received more than 1,000 comments in developing the core baseline over a year and a half.
- Pivot points and questions in looking at a consumer customer rather than an enterprise customer:
    - One was that customers buy products that they pull off the shelf.
    - Focused guidance on the device because that's the new part in IoT.

        - Introduces difficulty of management of the risks because of the interface to the physical world. However, it goes beyond the device because there's a digital/software component to this physical device.
        - How do we draw a boundary defining the product?
    - First pivot was, we cannot just talk about the device, we need to talk about the entire product.
    - The second one was moving beyond speaking about capabilities.
        - There were two workshops and two iterations in trying to identify criteria that could be used to underpin a labeling scheme.
        - The first iteration of the first workshop, we realized that the more specific we tried to get, and the more we tried to respond to every instance and every edge case, our criteria were either very brittle or they were going to be reduced to something that was meaningless.
- **Mr. Groman –** Commented that from his perspective, the problem is that there are sensors in the device and those sensors are collecting all kinds of data. For example, Amazon ring doorbell videos that were then misused. We must have data security also.
    - **Ms. Megas –** Replied that it's in the baseline regarding securing the data in transit, at rest, and protecting it. We didn't focus on issues around processing the data and use of data which are more privacy focused.
    - **Mr. Groman –** Commented that we need a better understanding of what data this IoT device is collecting, creating, storing, and sharing, because if you don't understand that, you can't possibly do an effective risk assessment and then have the right controls.
    - **Ms. Megas –** Replied that the recommendation is that the manufacturer does a risk assessment as part of their process and that they document the results of that risk assessment.
        - Don't know if we specifically call out data, but data is absolutely identified.
        - Don't know if we get specific regarding documenting the types of data that you collect and looking at the risks associated with that data. That can be additional guidance that we issue on risk assessment.
        - We did, at one point, try to go down the path of developing something that could help manufacturers conduct a risk assessment. One challenge we ran into was manufacturers of products only have half the story. They don't always know when they're building a product where that's going to end up or how it's going to be used. They may not know how a customer is going to use the data that they're getting off that device.
        - We tried to balance the idea of the baseline with a risk assessment and some minimum cybersecurity that should be built in.
- **Ms. Flynn Goodwin –** Mentioned that with all the customer issues flying around DC right now, this feels like it's ripe for additional clarification.
- **Mr. Lipner -** You talked about the manufacturer doing the risk assessment, but the manufacturer can't do the risk assessment because it's the end user whose data is at risk.
    - The consumer is not going to do an elaborate risk assessment, but there's probably only a baseline of information that the manufacturer ought to be making available to the to the consumer.
    - If the consumer cares, which they may or may not, they can at least understand in user friendly terms, what's going on and what risks they're taking.
    - **Ms. Megas –** Asked if he thinks this is a cybersecurity risk?
    - **Mr. Groman –** Replied it's both. He works with smart cities and as we start deploying sensors everywhere.
        - On the manufacturer side, there needs to be a fundamental explanation of all the sensors that are in any given device and their capabilities.

    - On the implementation side, understanding all the capabilities and the sensors' collection and creation of data that takes place.
- **Ms. Megas –** Replied that one of the baselines is about items regarding consumer awareness and communicating to your consumer but it is more focused on cybersecurity.
    - She agreed on the data side, there might be more to be done.
    - On a tangent, we explored a couple of things that might address what you're talking about.
        - We talked about putting together a transparency framework and we published an essay on it as the IoT team.
        - It was geared towards providing manufacturers and other actors in the ecosystem a framework to look at the cybersecurity of IoT information and data and to analyze that information addressing:
            - With whom to share this information,
            - What is appropriate to share,
            - Considerations for how the consumer is going to make use of the information.
        - There didn't seem to be a tremendous amount of interest.
- **Mr. Groman –** Asked if NIST is going to work on that side. An example is a paper about smart meters (which is part of IoT) and the reason why smart meters had challenges in early adoption was not out of security but the realization that the data was being collected, analyzed, and the amount of information you could infer about a homeowner based on the energy and when it's being used turned out to be extraordinary.
- **Ms. Megas –** Commented that she wishes she could have him come in and speak to one of her other boards. There is an alternative Federal Advisory Board focused on IoT and she thinks his input to that would be valuable.
    - They are preparing a report for Congress with recommendations on actions that the US federal government should take to support IoT adoption.
    - Trust is one of the issues identified.
    - If we cannot trust the technology, the US will not be able to adopt IoT to its full potential.
- **Mr. Groman –** Added an example of the installation of IoT sensors or cameras in federal public housing.
    - The claim is to make it safer, secure, or to enhance access to buildings and don't realize that the actual demographic and community we're seeking to serve are concerned. We just put up sensors without any explanation of where the data is going or how it's going to be used. That log process in any implementation of IoT needs to be part of that process.
- **Ms. Megas –** Commented while she doesn't want to speak for Federal Advisory Committee, she is the current convener of the federal working group that takes the recommendations.
- **Mr. Groman –** Added that the sensors are now extraordinary in their specificity.
- **Ms. Megas –** Commented there is an entire section of recommendations around what can be done to support privacy and IoT. They're still finalizing the language around those recommendations.
    - There was an early recommendation that we need a data framework to address the broader issues because privacy is a key consideration.
    - Regarding proprietary information we were hearing anecdotal information that there is a concern around the types of information that the sensors used in, for example, precision agriculture can be used to infer information about their practices for farming that have been handed down generation over generation and how that could be repackaged and potentially sold to their competitor.
    - Then there's the national security concerns around data and where data goes. There was call for a data protection framework that would look at:

            - Considerations around enabling the sharing of data because we also understand that making data available, especially to adjacent technologies like AI, is going to be important to the country going forward.
            - How can we make sure that we can still share the data? How are we addressing issues around protecting proprietary information, private information?
    - **Mr. Groman –** Added that talking about the data from IoT, it also matters about who the user is. It's not just privacy and civil liberties affected by the federal government using IoT sensors to monitor locations.
        - Also, a lot of this data ends up in commercially available databases, and what about when foreign governments buy that dataset?
    - **Ms. Megas –** Commented that was part of the reason they talked about a data protection framework rather than just privacy. They wanted to make sure that it addressed national security and all these larger issues. This is a big issue for this report to Congress.
- **Ms. Fanti –** Asked if they are talking about IoT privacy labels?
    - **Ms. Megas –** Replied that there is a recommendation but they're still being drafted.
        - There have been suggestions regarding rolling it into the existing US Cyber Trust Mark that's being launched under the FCC but there isn't consensus. There was also a discussion around a labeling effort for it. All the draft and interim reports from the FACA are public.
        - She'd be happy to have the chairs of the FACA come and brief ISPAB.
    - **Mr. Gantman –** Added that he hoped to have some time to discuss how we can have a labeling scheme, which encourages competition that doesn't have an artificial ceiling.
        - Once you meet the baseline requirements, there's no way to distinguish yourself above everybody else that met the baseline requirements. So what's the point in trying? There's no way to differentiate if you're better than average.
    - **Ms. Megas –** Commented that they were asked to look at levels of assurance.
        - There were a couple of reasons they decided against it at this time.
            - One was because we were talking about an individual customer, not an enterprise or sophisticated customer. By putting levels of assurance out there and leaving it to the consumer to try to make a judgment on when do I, based on the risk associated with this device?
                - Do I want to pay more money? Do I want to pay less money? Does it need a three? Does it need a two? There are sophisticated enterprise organizations that are struggling with making those kinds of decisions and those tradeoffs.
                - We felt if there was a program owner with the resources to look at a device or look at a product and say that this product warrants a higher level than this generic baseline, that then it would be on the program owner to create a profile.
            - The second reason was we also wanted to make sure that we were not creating a situation where you had a couple of very large, very competent manufacturers, they have economies of scale. They build it for one and they can put it into every single product and build it at a four when maybe all it needs is a two.
                - This would create a barrier to entry and unfair competition, where you could have more innovative smaller companies that also want to bring something to market, but they can't compete.
- **Ms. Moussouris –** Commented that she thinks they might want to be more prescriptive on the vulnerability disclosure piece than they have been and to make sure that we have time to discuss what happens at end of life.

- **Ms. Megas –** Replied that they do have requirements for end of life, they're just not very prescriptive.
    - We recognized early on that the device is only part of the BIG IoT ecosystem.
        - The unique things about IoT devices is sometimes they belong to multiple systems. They're part of the vendor system. They're part of your system.
        - "Follow the data" may be a way to understand the system and the boundaries of the system and where the data from the device goes.
        - Thinking of putting out a companion document because, for the labeling scheme, we are allowing the manufacturer to define what their product is and where their product boundaries are.
            - No way to come up with a "one size fits all" definition for those products.
            - Put a statement into 8425 describing the typical parts of a product,
            - Not every product has a companion mobile app, but many do. That mobile app gets you access to the device and lets you control the device and get access to the data.
            - Most IoT devices have a cloud component to be considered or a back end that may or may not be using cloud technology.
            - We stipulate that all the existing FISMA guidance still applies to all of that.
            - Our guidance is focused on what do you need to know about the device and what the device supports as that intersection between different systems.

Vulnerability Management

- There's this layer in the middle for specifying how you structure your vulnerability management processes or vulnerability disclosure.
    - Recommended that the program owner look at and adopt standards that meet the intended cybersecurity outcomes for vulnerability disclosures. In the timeframe that we were given as well as the broad scope of IoT, it would be too brittle if we tried to get too specific in our guidance.
- **Mr. Gantman –** Clarified that the program owner is the FCC in this case?
    - **Ms. Megas –** Agreed. Yes, it is the FCC in this case.
- **Ms. Moussouris –** Commented, regarding vulnerability disclosures and handling processes:
    - There are two international standards on vulnerability disclosure and vulnerability handling processes and organizations usually only try to follow one of them. It's functionally impossible to do the one without following the other.
    - CISA issued its first binding operational directive to the rest of the federal agencies saying you must establish a vulnerability disclosure program. They essentially said, figure out a policy and publish that, and then just have somebody ready to fix the things.
        - That's skipping over an entire standard.
        - They have completely lost the plot on the internal processes.
    - ISO/IEC 29147 is vulnerability disclosure. ISO 30111 is vulnerability handling processes.
    - Ideally, you have established your vulnerability handling processes well before you ask the community to tell you about vulnerabilities. A vulnerability handling process is how you deal with vulnerabilities that you discover from any source, internal or external.
    - The term vulnerability disclosure program (VDP) is often mislabeled as vulnerability disclosure policy,
        - you just can't throw out a policy, you need the infrastructure to support the policy.
    - If you include the note to have a vulnerability disclosure program, make sure you include both of those standards.

- **Ms. Megas –** Replied that she thinks that is a very good point.
- **Mr. Lipner –** Asked if the Secure Software Development Framework (SSDF) one of the standards that's referenced?
    - **Ms. Megas –** Replied that our baseline, and the documentation that we asked the manufacturer to have, includes what sort of secure software development process do you use?
        - Developed that way before the SSDF was published.
        - Anticipate that, as they go through this process, somebody could propose that the SSDF is the way they will meet the requirement to document a secure software development process.
- **Mr. Groman –** Commented that, in the context of IoT, you're using the term "labeling" to refer to ways of helping consumers be more aware of what's going on.
    - A new approach could be having just in time active labels and notices to consumers that sensors are operating, for example:
        - Having a new operating system for Apple where, if your microphone is activated by an app, an orange light appears to indicate one of your sensors has been activated.
        - An indicator if my toaster is going to be speaking to my refrigerator and then dialing back to tell the company about the conversation.
    - Privacy implications are not just for the homeowner.
        - Guests and their data in the homeowner's house are also affected, whether they ever consented to that or not, because they didn't know that the homeowner had an Alexa on.
        - So the risk of IoT includes contemplating that it may not even be my house and suddenly, my name is in something.
    - **Ms. Megas –** Replied that when we developed NISTIR 8228, it was very enterprise focused.
        - In the appendix of NISTIR 8228, we published the very first security and privacy baseline for IoT devices.
        - A multi-association letter was sent to the White House asking for that appendix to be turned into its own baseline for IoT devices. The use of indicators was one of the mitigations we brought up in the appendix.
        - In 2018, when we started the work on the baseline, it was suggested that we really shouldn't try to address both a cybersecurity and privacy baseline, that we should only focus on a cybersecurity baseline. We didn't have the NIST privacy risk framework at that time.
        - There were too many open policy questions still on privacy.
    - **Ms. Flynn Goodwin –** Added that now the timing is for this blended privacy and security assessment or conversation, because NIST is going to have to deal with this for AI.
        - It's time for us to get over those ivory towers of privacy and cyber security. IoT is a wonderful testbed for combining those two in principle,
        - The Europeans are getting tangled up in their AI Act. They passed the act in June and they're struggling with implementation,
        - This is a real opportunity to do some testing now to get ahead of this problem for AI.
    - **Mr. St. Pierre –** Added that in the AI risk management framework, one of the key characteristics for trustworthiness is privacy. Is it privacy enhancing?
    - **Ms. Megas –** Commented, we do create silos between cybersecurity and privacy and between IoT and AI. AI might be the brain, but IoT are the eyes, ears, legs and arms. You can't decouple the two and talk about them as if they're two separate things. We are seeing more and more AI in IoT products.
        - CPSC had a workshop on this almost a year ago and what I took away from that is they said, "we just want to make sure that the product doesn't harm the customer." However, they said, we do need to understand when it's a cybersecurity issue. When it's an AI issue. When it's a

design issue. Because if a consumer is harmed, we must understand how it is going to be remediated, and we must be able to test that it has been remediated to be able to put the product back on the market.

  - **Mr. Groman –** Added that cybersecurity for IoT suggests you're trying to protect something. In IoT, what we're trying to protect is the data that's been collected. I'm securing my IoT device because I don't want everyone using the camera. We can't divorce the two because the device has a sensor that makes it so sensitive. That's what the cybersecurity is there to protect.
  - **Ms. Flynn Goodwin –** Added that it won't matter because, even if we have a standard here, the Europeans are just going to regulate it and we'll all have to snap to that baseline. The more we can get really crisp on the data, the less we must live by what the Europeans decide. That ultimately is what keeps happening in this space.
  - **Ms. Megas –** Responded that we do address the protection of data and idea of confidentiality of data, the traditional CIA work, especially when we pivot into the product view, our cybersecurity outcomes around protecting data extend to data collected from the device that is in the backend and data that is on the mobile app. The part that we don't address so much is how that data might be used by those that are authorized to collect the data. Are they allowed to package it and resell it, or are they allowed to turn it into a report?

Report to APNSA

- It's a public report. Made some recommendations when we delivered the report based on what we heard in the pilots that were conducted.
  - There needs to be a consistent label. While there were lots of good existing schemes that exist out there such as the Connectivity Standards Alliance. There was quite a marketplace already of IoT cybersecurity schemes.
  - Some of the feedback is from a consumer perspective, they can't get to smart fridges and try to discern between two different labels. There would have to be a consistent label.
    - Consumer education is critical. To look at a product that has a label or doesn't have a label, you must know what you're getting by purchasing the product that has the label to drive demand for it.
    - You also must make sure that the consumer knows, for example, there might be the capability to set a password, but the consumer needs to also be educated in how they set the password on the device.
    - Feedback was that this needed to be owned by the government because it was too big for any single organization to take on and advocate for that. There must be flexibility for a wide range of products and approaches.
  - Liability considerations around a voluntary label needed to be addressed. This is coming up in the board discussions in my IoT advisory board.
    - If this is a voluntary effort, there were concerns around why would a manufacturer sign up to put a label on something that is making certain claims about the product, especially if there isn't a demand for it, and if customers aren't clamoring for it?
    - They are taking on a liability by making certain claims about their products. Why do it?
  - CMU did some interesting research on consumers and how consumers view labels. It wasn't binary: no label or something with a label.
    - The study showed a product with no label, another with one star, and another with two stars.
    - The customer will infer that the product with no label is better than the product that has only one out of two.

- **Mr. Groman –** Gave the example of a device that has a microphone that is always on and recording. That seems like a material statement that you want to make and disclose. If you can shut it off and how? You're right, different kinds of IoT devices have different sensors and functionalities and data collection. It won't be one size fits all. If my TV has a camera that's filming me watching so that they know how many people are watching, that seems pretty material to tell the user they're being filmed.
- **Ms. Megas –** Replied that the NPRM is closed, but she thinks there's still opportunity to comment to the FCC. There were lots of comments.
  - There's opportunity in the implementation for the FCC to say, we have an outcome that says you should make the consumer aware about features on the device and of the product.
  - It could always be addressed in the implementation that this is one way you must make the consumer aware of the product.

- International came up quite a bit.
  - What the EU is doing with the CRA coming down the pipe.
    - There was a lot of call for both self-certification and third-party certification and that we should allow for both, because some of these IoT products have very short lifespans, some of them are very inexpensive. Having strict third-party only certification may not be the most effective, especially when you're talking about products that they might have a new product or new version of the product every three months, four months, six months, and it's a low cost product.

Report to Congress

- The NDAA, of 2021:
  - Directed the Department of Commerce to stand up both an IoT advisory board and a Federal Interagency Working Group on IoT.
    - The intent is to look at all the barriers that we have to IoT adoption, and what the federal government should be doing to enable IoT adoption or remove those barriers.
  - The report is due from the IoT advisory board in January 2024 and finalized in February.
  - The recommendations for actions that the US federal government should take will go to the federal working group. We will prepare responses to that and deliver a complete report to Congress six months later, the IoT caucus has been meeting two days, almost every month because of the very short timeframe to get this report done.
  - The role of privacy, security, and trust comes up often and takes up much time. Connectivity is a big one as well.
- **Mr. Groman –** Asked if there is someone there with the expertise in privacy, data protection and trust?
  - **Ms. Megas –** Replied Yes, we do. There are 18 members on the board. It covers everything from agriculture to health care. There's a lot of focus on the consumer, even though it wasn't necessarily called out in the legislation.
  - **Mr. Groman –** Asked connected medical devices?
  - **Ms. Megas –** Replied correct. That's why it was one year to deliver a report to Congress.
    - It's looking at everything from smart traffic technologies to the use of IoT and logistics, and the use of IoT and supply chain.
    - The scope is nationwide, documenting for the hill what the big social economic benefits are, even looking at issues around climate change, sustainability, and how IoT can be used in that area. Also what can be done to catalyze that adoption.

- **Mr. Gantman –** Asked, regarding the attesting entity, whether it's self-attesting or a third party, is there a provision who it must be?
    - **Ms. Megas –** Replied that's up to the program owner and the FCC and how they implement it.
- **The Chair** thanked Ms. Megas for her presentation.

Software Liability Overview and Issues for the USG

Melanie Teplinsky, Senior Fellow Tech, Law, and Security Program, American University, Washington College of Law

Introductions

- Thanked ISPAB for the invitation to speak and provided some background:
    - Was a math major in college in 1994, when NIST announced the Escrowed Encryption standard and sparked the crypto wars. That debate is responsible for her being here today. She reached out to NIST and Miles Smid was the head of the computer security division. She spent that summer working encryption policy issues for the computer security division,
    - Was on the advisory board for CrowdStrike.
    - Currently serves as a senior fellow at the Tech, Law, and Security Program at American University, where she also teaches advanced cybersecurity law.

Insecure Software

- The problem that we're trying to solve is that insecure software is the norm today.
    - That's a significant problem for consumers, economic, and national security, all of which are put at risk.
    - Examples of this problem:
        - We can start with MOVEit.
            - The biggest hack from CL0P, a ransomware syndicate out of Russia, executed a massive data theft campaign. It successfully exploited file transfer software and breached over 2,000 organizations and 65 million individual. expected to reap about $75 million for the Syndicate. This is the third time CL0P hit File Transfer software.
        - The Microsoft Exchange hacks.
            - Chinese state sponsored attack on Microsoft's Exchange Server software, that gave them email access to hundreds of servers and organizations across the government.
        - SolarWinds, and recently
        - The Cisco vulnerability in IOS XE, an unknown threat actor that's exploiting a vulnerability found in routers, network devices, and wireless controllers with potentially up to 80,000 affected devices.
    - These hacks must be understood in context. Our digital ecosystem has been built against a backdrop of intense geopolitical competition among the world's major powers.
        - We have Russia, an acute threat to our interests and values, we have China, which the Pentagon has called the pacing challenge of our time.
    - All this malicious cyber activity is affecting us not only in terms of consumers and privacy, but also in terms of our economic and our national security interests.
- The Biden Administration's National Cybersecurity Strategy.
    - The strategy seeks to position the US and its allies and its partners to build our digital ecosystem in such a way that it's inherently defensible. It's resilient, and it's aligned with our values.
    - Looking to modernize and preserve our innovation and our competitiveness.

  - The administration has very explicitly stated that it intends to reshape the laws that govern liability for software vulnerabilities and other risks created by software.

Market Failure

- The concern is that to date, software developers have been prioritizing innovation and time to market over safety and security. Market forces are rewarding folks who are hurrying to get their product out and add more features, rather than add security. When this happens, we end up with products that are shipped with known vulnerabilities, products that are shipped with insecure default configurations, integrating third party software that maybe has unvetted or unknown features, and the end users are suffering.
- The starting point for software liability is the idea that there's a market failure.
  - Companies don't have the needed market incentives to produce more secure software. There are two reasons for this:
    - One is externalities.
      - Exploitation of software vulnerabilities don't fall entirely on the software developer.
      - A negative externality is the cost of doing business that the developers don't consider because it falls onto others.
    - The second piece is the liability exemption.
      - Software is largely exempt right now, from the usual liability regimes that would apply products and services. If you start with contract law:
        - You open a piece of software and download it; you must click through a whole bunch of terms. You click through them, you say, yes, right. You've just entered a contract.
        - Those click through agreements have been enforced by courts.
        - Software sales sold within supply chain, are governed by end user licensing agreements, etc.
        - Courts have upheld terms and conditions that disclaim liability for software problems that result in injury or damage.
        - Contract law is not providing software users with meaningful remedies for harm caused by insecure code.
      - What about tort law?
        - Tort law isn't much better. Turns out that software's mostly excused from liability under tort law for a bunch of policy reasons. We're concerned if we make it too expensive, if there's too much liability, and it's too costly. We're worried we're going to hinder innovation.
        - Courts have tended to say that hackers, not software providers, are the ones who are responsible for breaches.
        - Tort law is also limited as a matter of legal doctrine. There's something called the economic loss rule in tort, and it generally bars tort liability for purely financial losses because lawyers like to think of financial losses as the domain of contract law.
- **Mr. Groman –** Asked, “What about the individual?” When there is a software breach, we have the national security interest, and economics, but also 100 million people's data has just been compromised, and they're going to be victims of identity theft. That's an externality. It's not the company that’s going to feel that pain. It's the individuals and those individuals didn't sign a contract, and yet it's still almost impossible for them to bring a lawsuit. Can we do negligence? Reckless? Does that work with tort?

  - **Ms. Teplinsky –** Replied these are all possibilities. That's not what the administration is proposing if I read it correctly, but certainly there is a range of legal options. You aren't limited to a tort remedy.
- **Mr. Groman -** Asked where do people fit? It's not just big companies. It's also my data you lost.
  - **Ms. Teplinsky –** Replied that it depends on the specifics. Right now, there are a whole bunch of different ways to go at it; you can go as a data breach issue, a security issue, or a contract issue. There's been a limited success with consumer protection state by state. It's very context specific. There's also a role for insurance at the consumer level.
- **Ms. Moussouris –** Commented that bringing up the FTC, they have successfully gotten settlements with manufacturers about insufficient security protections.
  - One, in 2016, with the ASUS routers out of Taiwan. They settled, and it was partially because friendly hackers had tried to warn them about these issues. I'm interested in understanding how those levers may evolve.
  - **Ms. Teplinsky –** Replied that the FCC has been enforcing in this space for over a dozen years and what we've seen is that it hasn't made a substantial difference in terms of the quality of software. What we see is most of the settlements are relatively small.
- **Mr. Gantman –** Commented that he's going to challenge the notion that this has been a market failure. A lot of devices today are remarkably more secure than I think most of us would have guessed could be possible 20 years ago. There has been a lot of change over the last few decades. Software has done a lot more.
  - There is a lot of variance. But what's our point of reference? What is the segment of the economy, of the technology, outside of software, that is more resilient to adversarial attack? That's our ideal example of what it could be, and I don't think there's anything that's more robust than software today.
- **Mr. Groman –** Added, in other words, when we talk about market failure, we're asking what are the risks that are being incurred from this failure and are they being felt at a massive systemic level across our economy, nation, and individuals. And it doesn't fall on the companies that are producing the products.
  - **Ms. Teplinsky –** Replied that your question is one that's been raised by a lot of folks. There is not a general agreement. We don't have a way yet to effectively measure how good our security is. Nor do we have a reference point. There's a sense of frustration that, having worked on these issues for a decade or two, we haven't closed the holes in our system fast enough. We're still seeing serious implications.
    - We don't have perfect data. Should we go after low hanging fruit and start to make progress? While we're doing that, let's put in place some metrics-based approaches. Those efforts go hand in hand with the idea that we will get some metrics to be able to make better decisions.
  - **Mr. Lipner –** Mentioned he would disagree with the prediction that we're going to get more data on which to make our decision. Measuring product security and cybersecurity is hard. You can measure effort and things that people do. When you get up to the way the hacker looks at it, and breaking the system, then somewhere along there you're in the unknown.
  - **Ms. Teplinsky –** Added that the adversary matters here. Using the analogy of a car thief.
    - There's one type of car thief who walks around, and checks doors, and if there's one that's open, he goes to the car and steals it. There's another kind of car thief that is professionalized. It really doesn't matter how good your lock is. He's getting into your car.
    - So it depends on whether you have an adversary that is an opportunistic adversary, or a persistent adversary who is dedicated to getting into your system.

Changing the Status Quo: US & EU Approaches

- The strategy is seeking to change the status quo. The law offers lots of avenues.
    - Regulatory law, which comes into play through the FTC.
    - Tort law.
        - A tort is an act that gives rise to an injury or harm to another. It's a wrong for which courts impose liability. There are several forms of tort liability:
            - Strict liability is when you hold a manufacturer responsible for the bad cybersecurity outcomes caused by the product, regardless of whether they're at fault.
            - Negligence. Under negligence, standard manufacturers are liable only when a bad outcome arises from their failure to comply with the standard of care. They breached their duty of care when they built their product.
            - Safe harbor model. Safe Harbor model protects manufacturers from liability, if they comply with specified requirements, like a defined set of secure development practices.
- **Mr. Groman –** Asked if there is something between strict liability and negligence?
    - **Ms. Teplinsky –** Replied that some people are thinking about recklessness in terms of inverse Safe Harbor, where there are just certain practices that are so obviously blatantly, ridiculously out of step with what we expect, that they should not allow you to have the protection. I think that captures the recklessness concept. It’s just said in a different way.

EU Proposals

- There are two important initiatives in Europe that we should talk about: one is the resilience act, and other one is the product liability directive.

Cyber Resilience Act (CRA)

- The Cyber Resilience Act was introduced back in September 2022. It passed some key committees this past summer in Europe. Likely to be agreed on in 2024, which would mean that its obligations would come into effect in 2025 or 2026.
    - The purpose is to establish cybersecurity requirements for devices and software that are marketed in the EU. If that occurs, it will affect our market as well, for obvious reasons. Much like the National Cybersecurity strategy, the Act seeks to shift the security burden on to software developers, as opposed to software users. The theory is that software developers know how to mitigate these issues better; how to mitigate the vulnerabilities, distribute patches, etc. and that it's easier to mitigate vulnerabilities at the source rather than requiring users to do that.
    - At its core, we're talking about legislation that puts obligations on software manufacturers. The requirements of the Bill include specific obligations that would depend on the criticality of the software. As we think about our own software liability structure, there's a non-critical class of software and a critical class of software.
        - If you're non-critical, you get one set of requirements. Those that are critical get another.
- CRA imposes four major categories of obligations on developers.
    - One has to do with risk assessments. Some of them are very common-sense requirements such as:
        - Deliver your software without any known exploitable vulnerabilities,
        - Provide security updates.
    - It requires documentation. A description of the software design and development, and vulnerability handling processes.
    - It also would require conformity assessments.

        - This is quite interesting because for non-critical products, what's required is that the developer can perform the conformity assessment themselves by ensuring that they meet and then attesting to the fact that they meet the requirements set forth.
        - For critical products the assessment will have to be done by a “notified body”, which is an independent audit auditor that the EU certifies.
    - Then there are also some vulnerability reporting provisions.
- The scope of the CRA is limited to the software manufacturers publishing code in the EU. It excludes open source only to the extent that open source doesn't have commercial activity around it. The set of requirements around open source is very controversial.

Product Liability Directive

- Product Liability Directive.
    - For decades, there has been a widely shared legal opinion that software on its own doesn't fit within the definition of a product under either European or American law.
    - Recently, there's been some renewed interest in applying product liability law to software and most folks equate product liability law in software to strict liability. But product liability law is broader than that. There are cases of negligence that can be included under that.
    - The European product liability directive is a strict liability regime for products. That means there is no need to show someone's at fault. To get compensation for damages caused by a defective product, you just need two things. You need a product, and you need a defect.
    - The EU is in the process of updating their directive, which is about 40 years old. The Commission is trying to bring it into the digital age.
        - They propose to expand the scope of their directive to include software within the definition of product. This would include operating systems, firmware, AI systems.
        - These software products, include software stored on a device, accessed through cloud technologies, or a component of a physical product. There are exceptions. The Product Liability Directive would exclude from liability free and open source software developed outside the course of commercial activities.
- **Ms. Fanti -** Asked for clarification. Trying to understand if there's an inconsistency between the reporting requirements and the liability directive. You have a stronger set of reporting requirements for critical software, but the liability applies across the board.
    - **Ms. Teplinsky –** Replied it's a stronger set of obligations regarding what the software provider must do. They’re not reporting requirements, but a set of obligations that are being imposed.
- **Mr. Groman –** Mentioned that strict liability strikes seems difficult. How do you prove harm, injury, or damages? What is your cap? Does it play out in lawsuits by the private sector or is it a lawsuit by a government entity?
- **Ms. Teplinsky –** Replied that she thinks it changed the caps.

National Cyber Strategy

- In the National Cybersecurity strategy. Three key elements to what the administration is proposing. The prevention of disclaiming liability, the standard of care and the safe harbor piece.
- The administration committed to working with Congress and the private sector to develop legislation to establish liability for software products. And they specifically said any such legislation should prevent manufacturers and software publishers who had market power from fully disclaiming liability by contract.

- Establish higher standards of care for software on specific high risk scenarios, and shape standards of care for secure development.
- They also said they were going to drive development of what they called an adaptable Safe Harbor Framework to shield from liabilities companies that were securely developing and maintaining their software products and services.
- The idea is they want to improve cybersecurity, penalize poor software and exempt from liability product producers who are following reasonable practices.

- Shortly after the initial release, a working group came together.
  - It's comprised mostly of legal academics who are working to explore different approaches to development of this kind of liability framework to provide inputs to Office of the National Cyber Director (ONCD).
  - There are a lot of questions and very few answers. We're working on the answer piece, but it's a tough problem. The issues are both horizontal in the sense that they're pure cybersecurity or software development issues, but they’re also technological, it cuts across sectors and disciplines like law and computer science. So these questions are not straightforward. There's no one person that is an expert on software liability.
  - The administration is committed to exploring approaches to trying to develop this framework and ensure that we can build a software liability framework that would work.
  - Planning a symposium for the first quarter of 2024 and our group is working to actively develop some ideas to contribute to this.

Standard of Care

- One of the main challenges is to create a standard of care.
  - The standard of care is a legal term of art. It's used to determine when someone is acting non-negligently. If you fail to meet the standard of care, you're negligent.
    - Just because there's a security breach doesn't mean that a company has liability. A company can do everything right, and still be hacked.
    - What we're looking to do is figure out when a vendor has acted in a way that they should be insulated from liability.
    - The administration has suggested that the standard of care would reflect a minimum set of requirements to improve our overall national cybersecurity posture.
      - Some of these requirements might include good software design practices, testing, and vulnerability patching.
      - The tremendous diversity of the software market makes it very challenging to adopt a one size fits all approach because there's not one single generally accepted standard for securing software that we should just go out and meet.
      - Asking things like:
        - What standard should you have to meet before you go to market,
        - What’s the duty of the vendor to provide security support for product after it goes to market?
        - How long should the vendor provide support?
    - One approach would be to say software developers have to act reasonably. To not define the standard in any more detail. We've seen this. If you go back to the early 2000s and the privacy rules that were put in place by Graham Leach Bliley at the time in the financial services context, the cybersecurity rules basically say act reasonably and the industry really liked that because there was a lot of flexibility afforded by the reasonableness standard.

    - It can be applied broadly to many different types of developers developing software for many different functions. This approach means it’s also very difficult to predict liability and assess your risk on the front end. That was the complaint, that industry sometimes brought to the FTC. If you want us to act reasonably, you must tell us what reasonably is.

- **Mr. Groman –** Commented that industry would say it’s too vague and then turn around and say it’s too prescriptive. We must move beyond that.
  - **Ms. Teplinsky -** So that's the problem we face here.
  - **Ms. Flynn Goodwin –** Added that this will be extremely hard to apply to the retroactive ecosystem. You're creating a future standard of care for a future eco system. We can't solve this looking backwards. We can set a new standard to go forward and just deal with the applications that are going to continue to exist.
  - **Ms. Teplinsky –** Her sense is that we have to go forward and we're looking at a long-term problem that we want to solve over the next 10 to 15 years.
  - **Mr. Groman -** Most laws are proscriptive. The law goes into effect on a specific date, and then you have an implementation date, and it only applies to activity from that date forward.
  - **Ms. Moussouris -** Mentioned that a lot of the vulnerabilities that are discovered were introduced years ago. So even if you're trying to make sure that this is going forward, setting standards, even the things where the vulnerabilities themselves haven't been discovered yet, because they were introduced in the code years ago.
  - **Mr. Gantman -** Then the question of reasonableness comes in. You must assess what was reasonable when the product was made and that's very difficult. Also, by the time the incident happens, that product could have been bought three years ago.
  - **Mr. Lipner -** The putative benefits are out two years, three years, five years, eight years after the effect of that.
  - **Mr. Groman -** you can't go back, it's going to apply to software that is sold and placed in commerce after an implementation date, which is even after the effective date. It's not going to go back and look at all operating systems unless there are ties to something to come.
  - **Ms. Teplinsky -** From a 60,000 foot view, it's a transition problem. We have that every time we put in place, a new set of goals. It's something we can work with and figure out what is the right way to do this? How do we help and encourage companies to transition?

Safe Harbor

- We're trying to raise the bar in software security and design a regime that's incentivizing software providers to improve the security of their products without crippling innovation. Innovation is what powers our economic and national security. Policymakers recognize that if the risk of liability gets too great, innovators may abandon the market.
- **Mr. Groman –** Commented that the way we talk about innovation makes it seem as if innovation is always inherently positive. What we want is beneficial innovation. We need policies in place that are incentivizing beneficial innovation, maximizing the benefits from technical innovation while minimizing harm.
- **Ms. Teplinsky –** We completely agree. This isn't just about imposing liability for liability’s sake, the whole point of this is there are things coming down the road that we know how to do:
  - We know we can use memory safe languages, and we know how to do things like quantum proofing to some degree. How do we encourage them?

- How do we get ourselves to the place we want to be in 10, 15, 20, or 30 years? We're taking some of the first steps to try to do this, but we won't get it right the first time. This problem is too complex to get right up front.
- Safe harbors are somewhat different.
    - Liability is a stick. It gets companies to do something because they're worried about liability.
    - Safe harbors can be used as a carrot. They can encourage companies to comply with a set of requirements, and in exchange for the promise of an exemption from some liability. You can change the shape of that exemption, make it larger or smaller.
    - Safe harbor offers a lot of flexibility because it can be used as an incentive for companies to proactively address a problem in a way that liability can't.

- Existing models for Safe Harbor laws.
    - There is a set of state laws on safe harbors for cybersecurity. They encourage companies to take voluntary action to improve cybersecurity and meet certain prescribed standards.
    - The specifics vary by state, but the basic idea is that companies are shielded from liability arising out of a data breach when they develop and maintain the cybersecurity program that conforms to industry standards.
    - In most cases, there's some reference to the NIST cybersecurity framework. Some have suggested that's a model we should follow. There are reasons to be concerned about that model.
- There are several other very specific laws in the copyright space; Lockout Safety Act and medical practice safe harbors are debated every year.
- What's important is to pull out the considerations right behind these approaches.
    - An effective safe harbor needs to provide certainty and predictability to regulated entities and really needs to be defined clearly. This allows developers to do their job knowing what they must do to get protection. That reduces the uncertainty over how a court or jury will find a standard after that.
- Looking to create incentives to adopt best practices and avoid worst ones.
    - We need an updating mechanism because technology changes on a dime. If you don't have an updating mechanism, your rules are going to be outdated in three seconds, unless they're very broad, but they can't be too broad because we need certainty. Some folks have asked that there be some kind of metrics based cybersecurity. It would be good if we could know whether, and to what extent, any of the proposed actions have been effective. It can be difficult to determine that.
- We also have concerns regarding competition.
    - We see companies support regulations that have high compliance costs because it protects their markets against smaller competitors. We need to think about compliance problems and who's going to monitor a company's compliance with Safe Harbor, there's a spectrum of possible approaches to ensure compliance from:
        - doing nothing at one end of the spectrum,
        - self-attestation across the solution, as a self-report their compliance with the standards on the document that is not verified by some kind of secondary source.
        - On the other end of the spectrum, there could be a lengthy, expensive, external certification process, by a third party or a government, which should be designed to ensure compliance.
- There are lots of arguments about which approach we should take, if we should take the same approach for all entities, and how we should craft those requirements.
- Regardless of the certification model, the scope of the assessments needs to be clear and the criteria for reassessments needs to be clear.
        - When does change in what you're doing trigger a reassessment?

    - Do you need to reassess at regular intervals, etc.
- Final point on safe harbors. There are some things we shouldn't accept, and we should be able to start writing that list. Whether and how we can use all these legal levers to put in place stronger incentives for longer term improvements to secure software development that would move us toward some of these longer term goals.

Discussion

- **Ms. Flynn Goodwin –** Commented that was outstanding and asked if she has a sense of timing, because the creation of a voluntary framework, and the inputs for ONCD, also run the risk of running the clock out before something can get done within this administration and then we're left with the EU baseline.
  - **Ms. Teplinsky –** Replied that she agrees. Right now, the timeline is very slow; this strategy came out in March, the implementation plan came out in July.

There's quite a bit of work and smart people putting their heads together to try and solve this problem, but the reality is this will likely not be done during this administration. That's unrealistic for something of this magnitude.

- **Mr. Groman –** Mentioned there have been a handful of cases recently where the federal agencies have tried to enforce contracts using The False Claims Act.
  - **Ms. Teplinsky –** Added that this is a case where she applauds the administration. CISA and the DoD together have made some important changes to what they're doing. CISA in part has said, we're going to use the power of the federal government's purchasing to try and effect change on cybersecurity. Similarly, DoD has requirements in place. It's understood that there are many small businesses that are not in compliance with some of those requirements. So there has been more of a concerted effort to not only enforce, in the DoD case, but also to help folks who are having trouble meeting, particularly, cybersecurity requirements.
  - They've stood up some offices for the sole purpose of outreach to SMEs and others in this space to try and help them along. We've seen it in NSA as well with the cyber cooperatives. There's been an effort, not just in this space, to try and use the legal levers and to bring folks along.

Open Source Security Request for Information Review National Cybersecurity Implementation Plan 4.1.2

Anjana Rajan, Assistant National Cyber Director for Technology Security, ONCD/EOP
Nasreen Djouini, Senior Policy Advisor for Technology Security, ONCD/EOP

Introductions

- Background - Anjana Rajan
  - The Assistant National Cyber Director for Technology Security at ONCD.
  - New to the government, a cryptographer by background.
- Purpose for today: Give an update on our portfolio and what ONCD is doing; what we've accomplished in the last year and hint at where we're going in the next year.
- ONCD is the second newest component in the Executive Office of the President.
  - The impetus for why ONCD was created was to serve as the principal advisor on cyber policy and strategy to the President. One of our purposes is to serve as his convening function to make

sure that the departments and agencies are saying the same things collectively. Our way of demonstrating that, was the release of the President's National Cybersecurity Strategy in March.

- In the strategy, we're calling for some fundamental shifts, and that one of the shifts is that we need to rebalance responsibility of whose job it is to defend cyberspace. We want to shift it to those who are most capable. Far too often the onus falls on the small business, the state, local government, the individual consumer to defend against nation state cyber-attacks, and that's not tenable. Our theory is that the people who are most capable are the technical experts themselves.
- The strategy argues that many of the technical foundations of cyberspace are inherently vulnerable and every time we build upon that shaky foundation, new vulnerabilities increase, and our collective risk becomes higher.
  - Strategic Objective 4.1 argues that we must work to secure those technical foundations of cyberspace to protect our national interests.
    - One that we call out in the implementation plan is the open-source software ecosystem.
    - The log4J vulnerability in December of 2021 was an inflection point for the government because we realized that to have a secure resilient cyberspace, we must have a secure and resilient open source ecosystem.
  - We often start by thinking about the risks but for open source, we think about “if one component breaks, what's that ripple effect across the infrastructure?”
    - We worry about the interdependencies that can cause long term damage and put our adversaries at an advantage that they can exploit, such as the open source library.
    - Despite these risks, there is a way of thinking about this problem in a more nuanced and optimistic way because our stance in the Biden administration is that open source is a force multiplier for good as well.
    - The same principles that can make open-source software vulnerable are also the same properties that make it grow. We can’t have innovation and not include open source in that conversation. We also talk about this from a philosophical perspective; open-source isn't just a technical attribute, it’s an ethos, a philosophy, as well.
      - This idea that software and data should be transparent, it should be meritocratic; it should be available. That plays a key role in driving our values and our beliefs in justice and democracy.
      - We see this right now when we're looking at what's happening in Ukraine; two years ago, when the war first started, there were stories coming out about the open-source, cryptography community coming together to defend against Russian cyber warfare.
      - The concept of open-source software being an enabler for creating an equal playing field for other social issues, in essence, open-source software is a critical tool to share power towards those who are stewards of democracy.

- Focus on five key problems/outcomes and our focus areas for 2023.
  - First problem/outcome: ensuring that the Biden administration was aligned on the importance of open source software.
    - Being clear, as a government, why open source was critical to defend.
    - To get ahead of this risk, ONCD formally established the Open Source Software Security Initiative, also known as the OS3I.
      - An interagency working group that we lead. That goal is to convene the interagency towards putting resources, ideas, and policy solutions toward securing the open source ecosystem. It has grown to include very luminary institutions like NIST, DARPA, NSF,

OMB, CISA and many others. We've been working to engage the ecosystem, including civil society, the private sector, and academia in this work.

- Second problem/outcome: we wanted to figure out how we make sure that, as regulators and lawmakers are thinking about this work, the White House and the Biden administration is working closely with them to understand what we're learning.
    - One of the key outcomes that ONCD must focus on is how do we get ahead of the threat? How do we eliminate entire classes of vulnerabilities at scale?
    - The first vulnerability that we have chosen to focus on is memory leakage.
        - This is not the only vulnerability.
        - When the fiscal year 2023 appropriations bill report came out, Congress called upon ONCD to study the issue of memory safety. Six months later, we went to Congress, and we briefed both the House and Senate Appropriations Committee to share what we learned.
- Third problem/outcome: how do we engage with the open-source ecosystem themselves?
    - The open-source ecosystem is uniquely complex. It’s highly decentralized, internet native, and global in nature. It transcends borders by nature. Had to think about an international strategy from day one.
    - Working closely to engage our partners across North America, Europe, and Asia, to share what we've been learning through ongoing dialogues about our work. Our flagship engagement with the ecosystem was this past August when we were at DEF CON. We partnered with OMB, CISA, NSF and DARPA to publish a request for information on open-source software and memory safety.
        - This is important because when we talk about the ethos of open source, it's important for us to not just consume, but to contribute.
        - We see this RFI not only to inform our learnings, but to consolidate the expertise across this global ecosystem as a formal White House Research product.
            - We have 60 responses so far and we anticipate getting many, many more in the hours before the deadline. The RFI closes November 8, 2023.
- Fourth problem/outcome: We wanted to make sure that agencies, particularly those that are supporting critical infrastructure, have a plan on how they're both contributing and consuming open source software.
    - CISA published the Open-Source Software Strategic Roadmap in September. It is an important building block because it shows other departments and agencies who reliably need to figure this out, how they can start building those playbooks for themselves in future.
- Fifth problem/outcome: we're starting to take the first steps to start investing in the open source software ecosystem.
    - In September, NSF released a “Dear Colleague” letter that invites proposals to advance knowledge on securing open-source software. This letter is seeking submissions from researchers targeting software engineering frameworks, unsafe legacy code, dependency management, trust and safety incentive and organizational structures in education and workforce development. It’s been an important step to recognize the unique investment that the government needs to make to support part of this cyber ecosystem.
- These five workstreams discussed highlight the complexity of the open-source software ecosystem. This is one of the hardest cybersecurity challenges to solve. Our goal in 2023 was simply to say that, despite this difficulty making progress is not impossible.
    - These first five objectives are just the beginning. This is a multiyear, multi-administration investment.

        - Looking ahead to 2024, ONCD is starting to think about that next phase of work. Kicking off the new year by releasing a technical note.
            - The National Cybersecurity strategy calls for a rebalancing responsibility. Putting the responsibility into the hands of those who are most capable, and it's the technical experts who are most capable.
            - The technical note is a response to that call, and it is speaking to those people we are calling upon. We need to speak in their language.
            - She would like to find time one on one with some of the ISPAB members to talk about this in more detail, off the record, to get your thoughts.
- **Ms. Flynn Goodwin –** Asked if there are any USG Best Practices already in existence for contributing code to OSS projects?
    - **Mr. Scholl –** Replied that there's a DoD open-source policy that provides clarity for DoD federal employees' participation in OSS projects. It is specific to DoD. That is all I'm aware of.
- **Ms. Moussouris –** Asked why was memory safety chosen as the target?
    - **Ms. Rajan –** Replied that when she joined the White House a year ago, open source was the very first thing that came on her desk. It was clear that we needed to start showing how good engineering practices made for good cybersecurity policy.
        - She had seen her peers in industry and society talk about this issue for almost a decade.
        - We looked at attacks between 2003 and 2014 that started with buffer overflow and memory leakage issues that had catastrophic impacts.
        - We proposed that ONCD take on a more technical policy solution that can help move this conversation faster.
        - We talked about memory safety, with the a priori question of how do we eliminate vulnerabilities at scale in a preventative way?
        - The first vulnerability class that we were particularly interested in happened to be memory safety because there was a lot of it. We framed it as the first of several that we need to tackle.
        - Our answer was that we need to go out to the atomic units of cyberspace. One of those atomic units happens to be the programming language itself, because the language is fundamentally not secure and everything, we build in that language won't be either.
        - And that was how we initially framed it and within weeks, memory safety became a very popular term in the atmosphere. We were engaging with experts and other folks in the OSS community in the tech community. And we saw Google put out their report that they had migrated components of Android to Rust and that their own research showed that 50% of the bugs in Android were eliminated.
        - We saw research from another public research that showed you could see up to 65 or 70% Vulnerability reduction. This is a great example of Secure by design.
        - As the months went on, she started questioning if the message is getting lost when we talk about memory safety, because that's not actually what we care about. We don't care about the language; we care about the outcome.
        - Second topic is space cybersecurity. This conversation started with an attack on Viasat's satellites. So, our office found that space cybersecurity was another really important area to focus on. In March of 2023, right after we launched the strategy, we hosted the first big cybersecurity executive forum at the White House. We brought in CEOs from about a dozen of the top space companies into a SCIF.
            - We had the CIA, the NSA, and the NRO. They were briefed on the real, not hypothetical threats and then asked how do we move this forward? We agreed to lead a month-long

cross-country roadshow and meet with the space industry. We traveled to Los Angeles NASA JPL, Kennedy Space Center, and Cape Canaveral.

- We had an event at the White House again, went to Peterson Space Force Base in Colorado Springs and ended at the Space Center in Houston. Held five, half-day, technical workshops with staff level aerospace engineers on how we implement space policy Directive 5.
- Space Policy Directive 5 said we've got to secure our space systems.
  - Our hope was to add more technical rigor. We have these workshops with the space community. We talked about secure-by-design and memory safety came up.
    - The computer techs in the room were fully bought in. The aerospace engineers said, “You have no idea what you're talking about.” They said, you guys in the cyber world think that everything just exists in zeros and ones, but you're not thinking about the atoms and bits that we must deal with when you're dealing with that cyber physical nexus, critical infrastructure, and critical components that are not even on this planet. The constraints are different.
    - We went back to our original question which wasn't about memory, safe programming languages, it was about eliminating a class of vulnerabilities at scale. We asked, what makes for a memory safe language?
      - It has a garbage collector function in it that allows you to automatically clean up the memory that's being used.
      - What are the properties of programming languages needed in critical space systems in orbit? Right now, their belief is that the only language that meets those properties must be C and C++. So here we have a problem. We suggested that we broaden what we mean by atomic units. Maybe by atomic units we mean programming layers but maybe we also mean the microchip.
        - We look at the microchip and ask about eliminating vulnerabilities in that arena? Maybe we think about the compiler so that, if you have to use a language as opposed to the kernel, you could still create checks before you publish your code.
- Over the last few months our thinking has evolved based on talking to technical experts. A lot of what we're talking about is demonstrating that we must go back to, what is the objective function here, what are the constraints, and what mission are we optimizing on?
- When folks in the space industry say “mission”, it's more than cybersecurity.
-

Discussion

- **Ms. Moussouris –** Mentioned that she used to work at the Department of Aeronautics and Astronautics at MIT and she’s familiar with rocket scientists saying, “you don't know what you're talking about.”
  - One of the things she remembers learning from Mr. Lipner’s past experiences is that you could build a secure operating system that was secure by design, but they would not come. The adoption of these systems was difficult because you couldn't get the app developers to write apps that would be good. She’s worried about this big effort being put in in that area, if we've already learned a similar lesson.
  - **Mr. Lipner –** Added that was a different century. NSA used to have something called the Orange Book, which was a set of requirements for building security features and assurance into operating systems. A team he led had tried to build what was called an A-1 system, which was in large part

formally verified. What we found after we had spent $20 million, was that nobody wanted a system that secure.

- The risk is getting ahead of your ostensible market. You balance what you ought to do with what you think you can do.
- He loves the idea of memory safe languages where they can be used. I think you're right to go from an emphasis on memory safe languages to eliminating memory safety vulnerabilities. Because we're talking about hundreds of millions of lines of code that are decades old, that people are going to be relying on for decades to come. Maybe, in some cases, you can rewrite that in Rust, or C sharp or Java or something. In a lot of cases you can't.
- On a different soapbox, a guy named Matt Miller published a BlueHat Israel talk from 2018. He talks about the reduction in memory safety vulnerabilities across Microsoft code base. They did that mostly with developer training plus static analysis tools.
- That's blocking and tackling, whereas memory safe language would be a 95-yard pass. But the blocking and tackling can pay off if you do it.
- Brings another point, which is that the developer training and developer motivation is key and, in a lot of communities, missing. There was a study out of the Harvard Business School of why developers work on open-source and what their attitude is toward security. It got a very negative sense of developers understanding of how to do, and their commitment toward, security in the open source community.
- If you work for a company and people are developing as part of their jobs, you can make security part of the salary continuation plan among other things. A lot of those levers that a company has don't apply in the volunteer open-source community. That's a long-term culture change.
- In previous conversations on liability, you talked about changes that are going to take years or decades to come into effect. Changing the culture to where open-source developers care about contributing secure code is important. And it's not a quick fix.

- **Mr. Gantman –** Added that he somewhat worries that we're fighting the last war, because quite a few of us have spent decades of our lives trying to eliminate memory corruption vulnerabilities. One thing for us to consider is, even if we could snap our fingers and all memory corruption vulnerabilities have gone, there will be a few attacks that go away, and particularly a few targeted attacks. But, in terms of the overall scale of other attacks, does that change if memory corruption vulnerabilities go away?
  - **Ms. Rajan-** Replied that she has been very precise over the last year of how she talks about this. This is not the only vulnerability that we're solving. If you migrate to Rust, and you think about memory, safety, vulnerability elimination, you're not solving the entire ecosystem. The significance of memory safety and why we're taking it upon ourselves to tweak the language and be more precise, being the White House, in this upcoming report is, it's a shift in how you think, the shift in how we talk about policies when we say secure by design.
    - What we're really saying is, “how do we do preventative cybersecurity design?” What we wanted to show is we can talk about an example of what this looks like in practice, to illustrate that when we say secure by design, it fundamentally goes back to architecture decisions a CTO is making.
    - What we're trying to illustrate is that we must start thinking about how we talk about the affirmative vision, the market driven reasons to do this. We want to make sure that we're using it as an illustrative example because this is not a new discussion that we created, but that transition wasn't happening in Washington.

  - **Mr. Venables –** Added, if we're always conceptualizing that memory safety removes a whole class of things, strong forms of authentication eliminates a whole class of things, software supply chain control can remove another class…
    - He agrees with her point, to put the nuance around memory safety. It's not just, "everybody should kill C++, go learn more Rust", because even for the large tech companies, most of us are just figuring out how to get a percentage of the codebase secured.
    - You've hit on exactly the right nuance, that the interconnection between the hardware, the kernel, the compilers, and the software, even just with flipping on memory tagging extensions, doesn't require a lot of change to code. There is a significant benefit in getting people out of the "memory safety equals recode" to that architectural mindset point.
  - **Ms. Rajan –** Replied that her hope in this paper is that we don't even say the word C++, or Rust in it, and really focus on the broader themes.
  - **Mr. Venables –** Mentioned that you can probably appeal to the reliability concerns. A big part of failures has to do with memory leaks, which cause some kind of memory tagging issues. What we found is some of the resistance in the real time systems in upgrading for security, you can pretty much get them there by the notion of upgrading for silent failures.
  - **Ms. Rajan –** Commented that it's something that came out of the very first industry workshop we had at Lawrence Livermore. One of the key recommendations was how we get the infrastructure providers, like the GitHubs and the IDEs and all those folks, building development infrastructure to be part of the solution. Our hope is to show that the White House is listening and bringing technical rigor to a hard problem. It's great that we're even having this nascent conversation. How do we avoid an oversimplification of an important and complex idea?
- **Ms. Moussouris –** Commented that when she looks at the most common vulnerabilities that are reported through full disclosure programs, what we see are injection flaws.
  - One thing that you might consider in the future is looking at input validation as a major area, because the injection flaw game is still running and it's easier in a lot of ways to tackle that and the ROI on what you get if you can deal with those problems is pretty good.
  - The other major area is your good old fashioned configuration or outdated patches.
  - These are the things that we see reported most often in these programs.
  - Hardly any of the general population of hackers are really focusing on memory safety bugs. It's even a small percentage from the ones that were reporting "en masse" to Microsoft.
  - Question: how are you thinking about dependency stuff?
  - **Ms. Rajan –** Replied this is helpful for in thinking about what's next, what's the conclusion in this paper we write.
    - The nature of cybersecurity threats is that it's like this cat and mouse game. Once the defender understands the risk and can mitigate it, then the attacker finds newer places.
    - This goes back to what we're thinking. Rather than prescriptive, we're watching, observing and slightly cringing, at the same time around, are we getting are we educating? Under simplifying? And how do we bring that nuance? What's the next thing that we're hearing from the community?
    - We take this idea of prevention and outsmarting the threat, to figure out what would be the way to go after the next big surface area?
  - **Ms. Moussouris - S**ometimes the injection flaw or the input failure of input validation is the vector to get to the memory corruption.
- **Mr. Lipner –** Commented that secure by design is a great objective, and more progress on memory safety is a good thing. Security by design is a harder problem because it depends on exactly what your

systems are doing, and general guidance is not super helpful. Enabling developers to understand how to follow that guidance would probably be closer to fighting the next war. It's not so easy.

- **Ms. Rajan –** Replied that we wholeheartedly agree. That's why we're seeing the need to say, how do we look ahead and think about the next step. We do not want to have a checklist of things to do, and brick by brick we've secured it. That's not how this field works and so much of the way we think about this is in this more game, theoretical way; what are the smartest and fastest and most scalable ways that we can outsmart an ever-evolving threat landscape. That requires a framework and a way of thinking rather than a prescriptive or an aspirational goal. We need to get folks to think about the practicality of this.

Day Review and Meeting Recessed

- **Mr. Lipner –** Called for feedback and thoughts from the Board regarding letters or guidance and asked whether that discussion should wait until the end of the meeting on Thursday.
  - No proposals were made.
  - No public comment requests were made.
- **Mr. Scholl** – Things that were raised:
  - More resources on NVD
  - An issue around the source data and the communication of the known exploitable vulnerabilities (KEVs). Is it just Feds? Is it at scale? How does it get back to vendors?
  - There was a lot of discussion about IoT, IoT data, and IoT privacy.

Liability Workshop and Discussion on Liability

- **Mr. Lipner** – Brought up an academic workshop on liability in November at one of the Washington area law schools. He will investigate getting interested folks invited to it.
  - **Mr. Scholl –** Mentioned that it's his understanding that the leads from ONCD on software liability are also participating in that Workshop.
- **Mr. Lipner -** We might come back to a real discussion on liability, maybe with the ONCD people, but if not, we'll work with the after action in the workshop at our next meeting, which hopefully will be before the government's meeting in spring.
- **Mr. Scholl -** Yeah, absolutely. I talked to James Halpert and Paul Tao at ONCD who are running this, and they said they think they would be ready to come to the March 2024 meeting.

Open-Source

- **Mr. Lipner -** Does anybody want to do anything about open source.
- **Mr. Venables -** Every company I know has sent in a response to that RFI.
- **Mr. Lipner –** Yes, SAFECode met the old deadline. Okay, so I'm going to send people a note. The deadlines have been extended.
- **Ms. Moussouris -** I will say that a lot of times when these open source projects add resources, they add developers, they don't add security.
- **Mr. Venables -** Some of our commentary has been to encourage them to form a closer partnership between the government and the open-source security foundation (SSF) to be able to drive the funds that the Open SSF must have to connect security people with the package maintainers to do exactly that. I think a big part is giving them tools just like they do in our enterprises to enable them to do it.

- **Ms. Moussouris -** This is what we heard from open-source security foundation, when CSRD heard about log4J. They had the tools, they made them available, but log4J developers didn't use them. So…
- **Mr. Venables -** the big difference now with the outcome of your work is it's been forced into them a little bit more. But the thing that I'm still concerned about is of the 16 or 17, critical infrastructure sectors, very few of the sector risk management agencies are driving the critical infrastructure companies to partner with the Open SSF. Others sub sectors, like finance, have done some good work. The financial sector open-source thing is partnering with the open SSF sector. We're trying to encourage other sectors to do the same thing. Without the sector risk management agencies driving that, the White House may not be having as much effect as they could.
- **Ms. Moussouris -** If the drive would be centered around do you take input? If you do, now's your time. You must do a security review, or you must do something, because if that had been a directive and the log4J developers had to take it, they would potentially have avoided some problems.

The Chair adjourned the meeting at 4:15 P.M. ET.

Thursday, October 26, 2023

The Chair opened the meeting at 10:00 A.M. ET and welcomed everyone to the call.

A Briefing on the EU Cyber Resilience Act

Christiane Kierketerp de Viron, Head of Unit for Cybersecurity and Digital Privacy Policy, European Commission

Introductions

- **Mr. Scholl –** Welcomed everyone and thanked them for joining us.
- **Mr. Lipner –** Introduced the day's focus and that in this session we will be talking about the cyber resilience act. We're interested in the EU's perspective on liability.
- **Ms. Kierketerp de Viron –** Commented that it's excellent that we're able to connect like this and have an exchange on the Cyber Resilience Act (CRA).
  - Will cover the main principles of the CRA and explain where we are in the work.
  - There are three different strands of work in the CRA.
    - The first: Maybe the most important, currently negotiating the proposal.
    - The second: is on the standardization work we're carrying out to implement this Act.
    - The third: international discussions. We have international partners; the US is extremely important for us.
  - Overview of our proposal that we came up with it in September last year.
    - A bit more than a year down the road in terms of negotiations with the member states and with our parliament. Normally, it takes around two to three years to get negotiated. We're moving ahead at good speed. We hope to be able to conclude very soon, in any case before the parliament goes into recess next year, but that also means that we will need a political agreement, hopefully sometime in December 2023.
    - Negotiations are between what the commission has proposed, the member states who have their opinion on what it should look like, and our parliament will have their opinion of what it should look like. Together the three of us will establish the final text. Overall, we have clarity on a lot of things already.

- Main elements. The CRA addresses digital, connected products' hardware and software to be placed on the EU market.
    - There are some exceptions, like medical devices or cars, because they're already covered by other regulations that provide cybersecurity rules. Our regulation is based on what we call classic product legislation.
        - That means that, through our regulations, we put obligations on manufacturers or distributors, and partners. We establish some essential requirements for cyber security that must be met across the lifecycle.
        - We establish a risk-based approach, which means that manufacturers would need to consider the risks and consider these essential requirements against a risk assessment.
        - Where we deviate from classic product legislation is that we establish that the manufacturer must provide security of the product even after it's put on the market because it will look at the issues with hardware and software. A lot of vulnerabilities, you only discover as you go along.
    - Reporting obligations in the CRA. If you have an actively exploited vulnerability, you must notify the authorities.
    - Conformity assessment procedures.
        - A very large share of the products will go through self-assessment, but the critical products would either be obliged to follow a standard or go through third-party conformity assessment. Some products are considered so critical that they always must go through third-party conformity assessment.
    - Harmonized standards:
        - To facilitate the implementation, we're also working on developing harmonized standards. In Europe, standard development is done with the European standardization organizations. We have started preparing for standardization because we know that it takes time to get standards ready.
    - Market surveillance and enforcement actions.
        - Part of our product installation framework and the outcome is that all products would get the CE marking. A lot of products already have the CE marking to show that they are conformed with European rules. This is also for consumers to be aware.
- There are many more fine details to the CRA, but this is the basic framework. We are preparing some draft standardization requests and we have carried out a very large mapping exercise of available standards in the field, including international standards. We are holding them up against our essential requirements, to determine which are the standards that can give us a presumption of conformity.
- Working with the US EU-US. Summit, which took place six days ago between President Biden and President von der Leyen. There was an announcement that the EU will work closely with our American partners towards mutual recognition between the US IoT labeling scheme and the CRA and will come forward with a joint cyber safe product action plan, stipulating the work ahead.

Discussion

- **Ms. Flynn Goodwin –** Introduced herself and asked questions about the vulnerability disclosure pieces.
    - The regulation requires vulnerabilities to be reported to ENISA, which is more of a standards body than an incident response body, within 24 hours of discovery.
        - What happens once those vulnerabilities are provided into ENISA?

    - What happens if there is a disclosure of that information into the public domain?
      - When a vulnerability is found within 24 hours, companies don't yet have a mitigation for it. It's a very time-sensitive at that point, you don't want that information out in public.
      - What happens if, under the control of the authorities, that information gets leaked?
      - What are the processes that are being contemplated to protect, remediate, or support companies whose vulnerabilities get disclosed?
  - **Ms. Kierketerp de Viron –** Replied Yes, this is one of the points that has been raised.
    - They have received a letter expressing these concerns and have also heard from US companies. We've talked with many companies. especially marketing companies, who have been raising constraints,
    - We are not asking for every vulnerability to be reported. What we are asking for is actively exploited vulnerabilities. It means that you already have someone inside your product, and this product is very likely to sit in a lot of sensitive areas.
    - There have been a lot of concerns about providing a blueprint for more malicious actors to be able to copy what the other bad guys are already doing. This is not our intent.
      - What's important is what kind of information are we asking for? We're asking, within 24 hours, for a flag basically. This will be part of how we prepare for the implementation of the CRA.
        - We are very conscious about asking for “need to know” information, so that such information cannot be replicated.
        - A lot of people are talking about how this as a public disclosure of information. Informing a cybersecurity agency is not the same as a public disclosure.
    - The key thing is that we ensure a high level of security around this process, we ensure that this information is reviewed and only shared where there's absolutely no cybersecurity risk. This is already in the Commission's proposal. When assessing the information, if it is a very sensitive piece of information, the manufacturer can also advise that, on cybersecurity grounds, this is not something that should be shared. However, there is legitimate interest for a cybersecurity authority to be aware when we have products that sit in extremely critical areas that are under active exploitation.
- **Ms. Flynn Goodwin –** Asked if she could unpack that threshold a little bit?
  - If there's one computer that's being attacked by very sophisticated nation state actors somewhere in the world, that's what ENISA needs to know right away to share with the National competent authorities in the 27 member states, correct? Is that the threshold?
  - Once there's evidence of an attack in the wild, that would seem like a national security conversation at a member state level, not something that then will be shared at 27 member states, where it's not in a classified channel. How do you then feel confident that that information stays controlled and protected? And is it the one machine that starts this? Is it 20 machines? Is it 100 machines? What's that tipping point?
  - **Ms. Kierketerp de Viron –** Replied that the key thing here is we're working here on an internal market basis. If a company has a product, for instance a software product, that product is sold in all 27 EU member states. We have one internal market, and these things are sold throughout the union. There is always union interest in this regard.
    - The question is, who has the need to know it? We're having these discussions with the member states and the parliament on how to design this and it is part of the negotiations that we're in right now. Cannot go into more details but there is EU interest, because if you have software with a vulnerability, it is likely that all National Cybersecurity agencies have an interest in it.

        - The second question is how do we secure this information when it's being shared?
            - Having discussions on how do we share this to secure communication channels? No one is interested in having one point which is easily accessible for malicious actors. The information, when it's shared with our national cyber authorities, needs to be shared onto the right security protocols on a need to know basis. It's also important to say that the amount of information we're asking is not information that would, even if it ended up in the wrong hands, necessarily instantly mean that it's exploitable again.
        - What kind of information are we asking for?
            - It’s not a report that would allow for another malicious actor to exploit this information. We would like to have a flag that there is something going on in a product. You also must distinguish between companies that are fully capable of handling such things themselves and smaller companies that would like to have assistance.
- **Mr. Gantman –** Commented that there is likely not going to be just one agency.
    - One of the concerns is this sets a precedent. If we must do this for Europe, we will have to do this for every other market in a very short time, which basically means informing every single major government around the world.
        - Even if they just stick with one flag, that also means assuming that it was a nation state attacker and that means telling that nation state attacker that they were caught, right?
        - Other complexities. Within 24 hours, vendors won’t know which products are affected. To do the root cause analysis and figure out what other products may be impacted and whether they're relevant to the European market, takes way more than 24 hours. There's a concern that this: A) creates a disincentive to proactively looking for exploitation, and B) diverts resources from addressing the issue to dealing with reporting.
    - Regarding the threshold, do you care about all types of exploitation? For example, if somebody breaks digital rights management (DRM) and can stream high quality content, do we need to report that within 24 hours? If yes, that's going to generate a lot of noise.
    - **Ms. Kierketerp de Viron –** Regarding the last question, what would be reportable. We’re covering a very broad scope. We're covering many different aspects of consumer products.
            - We want to ensure that for a vulnerability that is being exploited, there is work to patch it and to share that patch, so that companies are doing the right thing.
            - Once you discovered that there is a vulnerability in your product that you're selling in Europe, and that this vulnerability is being exploited, then it should be reported.
- **Ms. Moussouris -** Introduced herself as one of the signatories on the letter mentioned earlier. She is also the co-author and co-editor of the international standards on vulnerability disclosure and vulnerability handling processes.
    - The idea of the level of detail not being enough to foster further exploitation needs to be explored further by organizations because it is often just a matter of saying which product has a vulnerability in it that is enough for bad actors to start looking and start being able to exploit things. This was the case when we had the WannaCry worm.
        - It was based on some vulnerabilities in SMBv1. All you had to know was there's something going on with SMBv1 and a bad actor would look at that and immediately find several exploitable issues. So, it is not correct to say that simply flagging software in which there's something going on is a safe piece of information to give out without the concern of something like WannaCry happening as a result.
    - The second thing is; we know you're doing a mapping of standards, which is good to prove conformity. However, the ISO standards for vulnerability disclosure and vulnerability handling

processes specifically prohibit telling any organization that is not directly involved in creating the patch or mitigation. There's a reason that is in those standards.

    - Those international standards were contributed to and vetted by over 100 nations around the world, including all the European countries and experts in all the European countries as well. All the experts understood that this information is volatile and is easy to exploit once you have just a clue about it.
  - What do you think the effects will be from EU requiring this in that market? Because, if the warning is shared with EU, even at a flag level, that warning is going to be required in every single other government where all these products are in those markets.
  - What will be the net effect to the long-term cybersecurity of the users in European markets given that it will instantly have to be shared with every other market and government around the world?
  - **Ms. Kierketerp de Viron –** Thanked her for the information and the letter. We have certainly heard your concerns, and these are also concerns that we're taking into considerations in our discussions on how to design a process and a way of working with this kind of notification in the most secure way.
    - We have a lot of discussions with our national entities about how this can be done in the most secure way. They are aware of the standards and what's in it but when they talk to a lot of companies and ask them what to do if you have something very bad, they go to a national authority that they trust and with whom they have close cooperation. It's not something new that there is cooperation with cybersecurity authority in the cases of these exploited vulnerabilities.
    - For concerns of the regulation's impact globally, we're in a situation where the information about the vulnerability is already in the hands of a malicious actor and we believe that, to protect our infrastructure and protect our key players, this information is also essential to go to those whose job is to protect their critical infrastructure.
- **Ms. Moussouris –** Asked if they have been thinking about making an explicit exemption for independent security research?
  - I know you're saying malicious exploitation only; however, often within that 24-hour period, there is a longer investigation that's required from an organization to know whether it's a friendly researcher trying to exploit something versus a true malicious actor.
  - Explicitly excluding vulnerability research would add some clarity and some wiggle room for organizations to make that determination and it may take a little bit longer than 24 hours from the first time that they've witnessed the exploitation attempt.
  - **Ms. Kierketerp de Viron –** Replied that's an important point and something that we're very keen on clarifying. We say you have 24 hours to notify activity "by a malicious actor". It is once you have reason to believe that it is, indeed, a malicious actor and not a security researcher that the clock starts ticking and you notify.
    - Security researchers are something that we're looking at very carefully from the EU side, because we're also aware that, in several jurisdictions, there's no safe harbor; there's no distinction between when you are an ethical hacker or a security researcher or not.
    - How do we protect these people who are doing good things? This is a very clear thing with the CRA. The focus is only on the malicious actor. We do not want to get in between the very good work that security researchers are doing. In the letter there was some concern about how it affects the relationship between a security researcher and the manufacturer. The way the CRA has been designed, and the way it will be implemented, there's absolutely nothing that should affect that negatively. This is very important for us.

- **Ms. Flynn Goodwin –** Asked if they see any scenarios where there could be a national security exemption, for example: if there's a report to the national competent authority and a member state can notify the other members that it has received a report and that it will notify the others within a certain amount of time. Do you see a member state being able to become the owner of that, as opposed to ENISA, from a national security perspective**?**
- **Ms. Kierketerp de Viron –** Replied that we're discussing and negotiating who should receive the notification.
  - When you look at the position of the member states, they would like for their national CERTs to receive the notification. If you look at the position of the European Parliament, they would like ENISA to receive this notification. The setup is still very much under discussion and negotiation.
  - The other question was on the timing for ENISA or a CERT to pass on the information. It's important that we establish the rules for when you can withhold information and when do you pass them on? There's an additional level of security built into the process and that would apply whether it's a CERT or ENISA.
- **Mr. Lipner -** Regarding the certification practices that vendors are required to follow, 2 concerns:
  - The cost of doing the certification, whether third-party or self-attestation.
  - The effectiveness of that regime. If the set of practices is necessary and effective to reduce product vulnerability, then he's happy to see organizations spend money doing that.
    - Our experience with the correlation between certifications and resistance to attack has not been especially good, particularly in the software world, going back to the days of the US Orange Book in the 1980s and common criteria. Those standards tended to drive the introduction of security features effectively, but they haven't had any effect on resistance to attack, particularly in software.
  - How do you assure that the certification and regime practice are going to be effective in enabling you to be confident that you're getting secure software?
  - **Ms. Kierketerp de Viron –** Replied that either way (self-attestation or third party) would bring costs upon our industry. At the same time, this is always how we do product legislation in the EU, so, to a very large extent it's not a new concept. Our industry is used to this, and this is one of the things that industry specifically asked for. (Lost audio on the call.)
    - That's why the standards are so key; the standards will also help our businesses selling in Europe to comply. One of the key things to ask yourself is would it help?
      - We have had very good experience with this kind of framework, because we also have the market surveillance authorities, and we have the penalties in place as well.
      - Not advocating to use them anytime soon but it's more than the classic voluntary certification scheme. This is really a framework that, if you're not complying, your product will be pulled from the European market.
      - The level of incentive to comply with the rules applied through security by design, the appropriate testing and scanning, and the appropriate patching, is high.
      - What is at stake is that you lose access to the market if you're not doing things appropriately.
  - **Mr. Lipner –** Commented that it's clear, if vendors must certify to sell in the EU, then they'll certify to sell in the EU. That's not an issue.
    - The concern is that we don't have a good track record of software certification regimes resulting in more secure software.
      - The US has effectively mandated suppliers to the US government go through a regime defined by the NIST Secure Software Development Framework. There's not yet a formal

certification regime against that and we don't have a lot of experience with it to know what its effect will be.

      - The history of the software certification world, going back to the Orange Book, has been a lot of introduction of requirements before we knew whether they work or not.
      - Understand that EU has a long history of product certifications, but software certification is a different animal.
    - Recommendation would be that you get some experience. If we require software developers to do a list of things, does that result in improved software security or does it just cost them additional money without the benefits we're seeking? That's really the core issue.
  - **Ms. Kierketerp de Viron –** Replied that this is very important, because this is also the first time that we really regulate software security.
    - The CRA requirements are not very specific requirements. They're objective driven.
      - For example: we're asking to ensure appropriate authentication. Our list of essential requirements are very high level, the basics of security by design.
      - We have consensus that they are good and do have an impact on the security of the products, specifically of software. What we do say is that companies must do a risk-based approach, they must look at the end product, and then they must look at the essential requirements and the extent to which they would apply to them, they would have to find the appropriate levels.
      - The standards are what is forming these high-level updates into something tangible for the industry. Very few companies have raised any concerns with our CENSA requirements. Overall, those who are doing a good job at secure products, say these are things they're doing anyway.
      - This is very much a question of getting those who don't so actively consider security when they're developing their products to take these things into consideration.
- **Ms. Moussouris –** Brought up the complexity introduced by the hardware and software supply chain, how difficult that coordination is, and how these requirements may affect supply chain issues in unintended ways.
  - Specifically, if we look at the Log4J incident, if any of the product manufacturers had noticed exploitation of their products before the maintainers got notified, that notification requirement would have further disorganized an already chaotic process.
  - She had started Microsoft's first multi-party vulnerability coordination program about 15 years ago and it's not gotten any less chaotic to do multi-party and supply chain vulnerability coordination.
- **Ms. Kierketerp de Viron –** Replied that we see this as a way of better organizing this process.
  - We're aiming to achieve, with the regulation, clarity about who's responsible for handling vulnerabilities and putting the responsibility with those who can deal with it. That's the logic of the CRA.

National Security Memo 10 and Post Quantum Cryptography National Cybersecurity Implementation Plan 4.3.1

Dustin Moody, Post-Quantum Cryptography (PQC) Project Team Lead, NIST
Dylan Presman, Director for Budget and Assessment ONCD/EOP

Introductions

- Thanked the Board for inviting him.

- Background.
  - Dylan Presman, the Director for Budget and Assessment at the Office of the National Cyber Director. ONCD is the youngest office within the Executive Office of the President.
- Will be talking about the opportunities that quantum technology represents, the risks that they represent, and then where we are as a federal government in that area.

Opportunities

- Quantum technologies are going to be amazing for our society, for America, for our economy, and for our nation. The revolution that quantum technologies represent have been characterized as, quantum computers will be as like current computers, as current computers are to the abacus.
- Three kinds of quantum technologies: quantum sensing, quantum communications, and quantum computing.
  - Quantum Computing is the one that people talk about most.
  - Quantum sensing is already being used in many areas, including mining, and it will get only more sensitive and revolutionary.
  - Quantum networking is an area that's going to be critical in the future.
- He described and elaborated on a presentation slide showing how different quantum computers are to regular computers. It’s a Venn diagram showing what regular computers can do versus what quantum computers will be able to do, and the unknowns.
  - Quantum computers will be able to do a lot that regular classical computers can't do. There's a quote from a Microsoft researcher that says, “it will take the strongest computer that we have on this planet 1 billion years to crack RSA encryption. It will take a quantum computer 100 seconds.”
  - Quantum computers, once they're sufficiently mature, will be able to crack all the forms of encryption that we currently use in our unclassified environments. Primarily RSA, Diffie Hellman, and a couple of other ellipse type algorithms.

Risks

- Three areas:
  - If we accept the premise that quantum computers, once sufficiently mature, will break all known forms of encryption in our unclassified arena, there's a couple of issues that spin off that.
    - One is that transitioning to a new form of encryption, new algorithms is a lengthy process, and we've never done this sort of transition at this kind of scale. We have had a couple of cases where algorithms have been broken and NIST has pushed the federal enterprise to get away from those algorithms, such as SHA-1 that was broken probably a decade ago or more. We still haven't fully transitioned away from those.
      - The transition that we're talking about here is significantly more all-encompassing, therefore, the transition timeline itself is going to be onerous.
      - It will be a significant burden for federal agencies, critical infrastructure, and industry in general.
    - Two: The other vector of this threat that needs to be considered is harvest now, decrypt later.
      - We already know that our adversaries are gathering up large troves of encrypted data and storing it because they know that eventually they will have a quantum computer capable of breaking that encryption. This means that we must think about the lifespan of that data when calculating when we need to be transitioning or thinking about when we need to be transitioned.

        - Classified information: we like to keep it safe for 25 to 50 years.
        - Personal data and medical data: probably want to keep that safe for the lifetime of the individual involved.
        - The name of an informant, names of spies, positions of spies probably want to keep that safe for a long time.
        - There's a lot of intellectual data, intellectual property that we want to keep safe for long periods of time. Some of it runs out quite quickly, but there is intellectual property that we want to keep protected for long periods of time.
    - Most experts think that a sufficiently mature quantum computer is probably still 20 years away, maybe a little bit sooner, maybe a little bit longer. We don't know exactly.
        - When you think about if we need to protect the data with an encryption method that can't be broken by a quantum computer, and the quantum computer is 20 years away, you can already see that we're too late for that.
        - There is an urgency caused by this harvest now decrypt later form of attacks that brings that timeframe forward. It won't be sufficient to wait until the computer turns up. By that time they will already have quite a lot of access to our information in a way that would be harmful to our society.

Key Considerations

- Legacy IT
    - Still have a lot of legacy IT in the federal government, particularly within critical infrastructure.
    - Lots of technology that's been around for 20-30 years, that's already too weak to handle current forms of encryption.
    - The transition to post-quantum cryptography will be that much more burdensome.
- We have three buckets of technologies:
    - Commercially available off the shelf technologies, software, and hardware; that's the Microsoft suite, the Apple iPhones, etc. We can assume that the big vendors will take care of their own encryption transitions, and that will be relatively seamless within the federal government. We'll need to keep an eye on that as not all the players in the commercially available software markets are large.
    - Government customized stuff.
        - We have a lot of government customized technologies within the federal government. That will be a matter of working with vendors to ensure the transition takes place.
    - Legacy IT.
        - This is where 80% of the work is really going to have to focus. These legacy technologies are fully embedded across the federal government.
- Another big critical issue within government is that we tend to do our budgets quite short term.
    - Budgets are created in one year time slots, and they're not thinking long term. A ten year transition is difficult for the government to get its head around.
    - Convincing Congress to pass next year is challenging, let alone what may be happening in the next ten years. That's something that will have a particular impact in the federal arena.
- Cryptographic agility
    - This may be the biggest or the most impactful cryptographic transition that we have been called on to make, but it certainly isn't going to be the last we can expect.

  - With AI moving as fast as it is and quantum technologies coming online, there's going to be more and more breaks of cryptography. It behooves us to say, we can see that coming, and we need to build cryptographic agility into our into our future.
- **Mr. Venables –** Regarding cryptographic agility:
  - That's a critical point. Dustin might talk later about the likelihood of some adjustments having to happen for the PQC algorithms, given the new field of mathematics that many are based on.
  - There's going to be other breakthroughs and what's not landing in the private sector, and maybe in many parts of government, is the fact that this isn't just a crypto swap out or an addition of that. This is a re-architecting of how they're employing cryptography to have that future agility.
  - There's a need to hammer that crypto agility point way more aggressively, especially given the timeframes to do a PQC migration. Are you thinking, from a White House perspective, about how to campaign in industry to help people get a grasp on the timeframes?
  - **Mr. Presman –** Replied that he agrees.
    - The Office of Management and Budget (OMB) has put out specific guidance on how to do the transition to post-quantum cryptography. In that guidance, they speak to the need for building cryptographic agility into future generations of technology. Not a lot of folks read OMB memos. We're not that great at the broader campaigns and articulations.
      - That's where industry is good and where NIST collaboration, through the National Cybersecurity Center of Excellence (NCCoE) and the collaboration with industry partners, is important in building that groundswell.
  - **Mr. Scholl –** Added that at every NCCoE project and every PQC workshop; crypto agility is brought up and publicized that way.
  - **Mr. Presman -** But it is a tough thing for the federal government to say things in words that are interesting to the public at large.
  - **Ms. Fitzgerald-McKay -** Is there any document or standard guidance for how to achieve crypto agility for enterprises?
  - **Mr. Scholl -** Not specific to crypto. We have guidance on algorithms and key transition and transitioning, and we have guidance on key management at enterprise down to product level. We probably could generate some specifics about what agility is and how it works.
    - One of the byproducts of the NCCoE workstream is, with the use cases of PQC transition, they also want to co-develop playbooks to be for any transition. This is just the use case.
    - One frequent question is how do we future proof? The answer for the future is agility.
- **Mr. Gantman -** A related question, but more on the protocol design than infrastructure rollout.
  - A lot of the effects of switching crypto algorithms are at the higher levels of the protocol.
  - Is there any kind of standard guidance for protocol designers on how to design protocols for crypto agility? Not crypto protocols, specifically, but general communication protocols or application protocols that have crypto in them?
  - **Mr. Moody -** NIST doesn't get involved as much with the protocol level. We do a little bit but a lot of that gets done in other bodies, like the IETF or other places. For PQC, they are working on that. Regarding crypto agility, I'm not too aware of anything specific.
  - **Mr. Venables –** Added he's concerned about three things:
    - One: industry is not processing what crypto agility means in terms of the actual work.
    - Two: at the end of the transition, everybody's not understanding the knock-on effects higher up the stack. A classic example, some bank payment systems have X. 509 certificate signature fields embedded that will require everybody to change software and databases. Like a Y2K style problem that there's going to be other software changes.

    - Three: Regarding crypto agility, it's pleasing to see some of the other standards communities, like 6G, have finally processed that they must think about PQC because, by the time 6G fully rolls out, we'll be in the era where we're close to a viable quantum computer. But we haven't hammered home crypto agility in the standards community.
    - Concerned that many in the standards community will declare victory that they're PQC ready based on the current standards, but they won't be crypto agility ready, and therefore we'll be building in that legacy.
  - Not seeing, across the government, a focused effort to get those three things across. Either NIST, ONCD, the NSC, DHS, or maybe even NSA cyber collaborations needs to take this on as we're missing that piece.
  - **Mr. Presman –** Commented that a lot of work is happening, but you wouldn't necessarily be seeing it. We're early in this process in the transition. Now, where we are in the transition is in the preparation for NIST standards being finalized. Once NIST standards are finalized, the federal government will come out with a plan of action. We're still doing inventories and figuring out what our priorities are.
  - **Mr. Venables –** Suggested that NIST, or whoever, should up the communications on this so we don't miss opportunities.
  - **Mr. Gantman –** Added that those standards are still being defined with how many bits do we need for that identification tag, and the only thing that's changing is the exact number of bits that are reserved. That's not really agility.
  - **Mr. Venables -** The agility is not just the ability to swap out the algorithm. The agility is for keys and signature sizes to flex all the way up and down the stack. Everybody's not really thinking about that degree of agility as well.
  - **Mr. Lipner -** At the risk of piling on, we knew about postquantum and tried to build in crypto agility when we wrote our secure development standards internally, back 10-15 years ago. We didn't know what that meant in terms of lengths, several protocol exchanges back and forth to get something established, state or no state. The government, NIST, and the NCCoE project's real service is by providing more actionable guidance, concrete examples, to help with that. Of course, we needed it maybe two years ago.
  - **Mr. Presman -** Agreed.

Quantum Policy Timeline

- NIST took the lead in 2016, when it went out to the community and said, "Let's start thinking about what quantum resistant algorithms would look like," and has been pushing forward ever since.
- Congress passed the National Quantum Initiative Act back in 2018. The US needs to be a world leader in that field.
- Part of the issue is Executive Order 14028: Improving the Nations Cybersecurity.
  - Encryption is central in zero trust architecture. One of the central planks of zero trust architecture is encrypting data at rest and data in transit.
  - If quantum computers are going to be able to break encryption that could invalidate our architecture.
- The big policy year was 2022.
  - There were some significant breakthroughs technologically in the quantum arena such that, by 2022, most people have come to the conclusion that sufficiently mature quantum computers inevitability is not a dream being chased anymore, it's an engineering challenge to be overcome.
  - There was more investment in quantum technologies in 2022, than all previous years put together.

  - National Security memorandum 10 (NSM 10)
    - Two facets: one is that quantum is great, and we need to be a leader. The other is, there's a threat and we need to be cognizant of the threat and preparing for that threat.
    - NSM 10 directs the federal government to transition to post-quantum cryptography and to finish that transition by 2035.
  - OMB followed up with a memorandum that gave departments and agencies specific guidance as well as referencing earlier M-23-02 specific guidance on how to begin that transition by developing inventories of their cryptographic systems, beginning to do assessments and cost estimates for the transition to post-quantum cryptography and beginning the prioritization process.
  - At the end of 2022, Congress passed the Quantum Computing Cybersecurity Preparedness Act, which codifies those policies into law.

Where are we?

- Federal agencies developed the first pass at doing an inventory of their cryptographic systems.
  - When we did our first inventory of just physical IT technology endpoints, we were factors of 10 wrong using CDM (CISA's flagship technology system), continuous diagnostic and mitigation, tool to do inventory.
  - Designed this process, understanding that the first time agencies do an inventory of their cryptographic systems, they're going to have a hard time,
  - First shot of those inventories came in in May. The first cost estimates came in in June.
  - In August, NIST released the draft standards for three of the four quantum resistant algorithms that were their finalists.

The Quantum Threat

- NIST is looking to replace our vulnerable algorithms, which are the public key cryptography standards. Quantum computers will impact the symmetric key algorithms as well, but not as drastically. We don't have to switch out the algorithm. We might adjust the key size.

NIST PQC Competition

- NIST ran a cryptographic competition with the community. It's been going on for several years and it has done very well.
- We were looking for digital signatures and public key encryption or equivalently, key establishment.
- Established at the beginning that we'd need more than one algorithm.
- It's a newer research field, so we want to have a diversity of security assumptions to base the algorithms on. We have different performance profiles, so we need to take all that into account and have a small number at the end.

First Three Rounds

- We went through a few rounds of evaluation and analysis. Put all the specs online.
- Internally, we looked at them; people around the world looked at them; a lot of them were broken, a lot of benchmarking, and a lot of experiments going on. At the end, we had seven finalist algorithms and eight alternates, based on lattices, codes, multivariate, several different mathematical properties.

Round Three Results

- For public key encryption, or a key establishment, we picked the Crystals Kyber algorithm.

- For signatures we selected Dilithium, Falcon and Sphincs+.
- We have a lengthy report that gives all the details as to why we picked what we picked.
- There's one KEM and three signatures. We want to have a few more KEM algorithms in there.
- In the fourth round, we have a few more KEM algorithms that we're still evaluating.

The Selected Algorithms

- Kyber, Dilithium, and Falcon are all based on structured lattices. They are larger than the current algorithms we use in terms of their key size and their signature size, but not too badly.
- We think most applications will be able to handle using Kyber and Dilithium. Those are the two main algorithms that we expect people to use for most of their applications.
- Falcon was also selected; it has smaller signatures than Dilithium, but it has an implementation that is way more complicated. Some applications will be able to handle the floating point operations that are needed and get the smaller signatures, but many applications won't want to do that.
- To not rely solely on lattices, we also selected Sphincs+ for signatures.
  - More conservative design based on hash-based cryptography. The tradeoff there is its performance is not as good as the lattice based algorithms. It's a backup algorithm in case you need the higher security and performance isn't a top concern.

Timeline

- Right now, we have the draft standards for three of the algorithms out. Public comments for them are due on November 22, 2023. We will then review all the comments to make any changes needed. We don't anticipate any major changes.
- We expect to publish the first standards in the first half of 2024.
- We'll have another workshop next year and we have lots of ongoing standardization work going on, but the standards are the focus that we're making sure to finish.

Standardization

- The first three standards are coming out as their own individual standards called FIPS, a Federal Information Processing Standard.
  - We gave them new acronyms to differentiate them from previous versions of these algorithms. Their numbers and their names are listed (see presentation slides).
- The fourth algorithm, Falcon, is currently being written. It'll be published in draft form probably at the end of 2024.
  - We will also update and publish other documents to go along with these standards documents that have timelines associated with them such as guidance on KEMS because Kyber is the first KEM that we're working on.
  - We'll update our other standards that go along with PQC to have more details and guidance.

KEMs in the 4th Round

- We still have a fourth round where we have a couple of KEMs under consideration.
  - None of these are based on lattices. When we select one or two of these, we will have more KEM algorithms with different security assumptions.
  - SIKE, the bottom one, was broken, so don't use it. The other three, Classic McEliece, BIKE, and HQC, we will pick one or two of them to standardize. That decision will probably be in a little less than a year.

An On-Ramp for Signatures

- The onramp signature process.
  - Sphincs+ is our backup, non-lattice algorithm, but its performance is not as good.
  - We would like to have a more general-purpose digital signature algorithm, not based on lattices with better performance.
- We issued a new call about a year ago to get new signature algorithms to consider them for standardization.
  - This will be a multi-year process.
  - The first submissions came in: 50 were sent in, 40 were accepted. They're currently in the first round of evaluation.
  - Any standards that come out of this will probably be four to six years in the future.

Stateful Hash-Based Signatures for Early Adoption

- Have some stateful hash-based signatures that are already standardized.
  - These started with some IETF standards.
  - They're getting a lot of interest right now because NSA allows them for code signing and firmware updates.
  - They're not general purpose because you have to manage the state, but if you can successfully do that, with the restrictions that are in the standards, they're appropriate for those use cases.

Other Standards Organizations

- A lot of the other standards organizations and groups were waiting for our NIST process to finish.
- Now that we've announced the algorithms and are standardizing, they're making their own recommendations based off the algorithms that were in our process.
- We're working with ISO to standardize Kyber, and the other signature algorithm as well, so that we have international standards.
- Other countries are also making recommendations and their own internal standards.
  - For the most part, they are using the algorithms that were in our NIST process and have been heavily scrutinized. Exceptions, like China and Russia, are doing their own thing, but most of the international partners that we regularly work with were very involved in the process.

Transition and Migration

- Our main goal is to transition as much as possible by 2035.
- NIST will have more guidance, including with things like hybrid paths, where you combine the classical and PQC algorithms. That's not going to be mandated or anything. That's going to be up to the applications and protocols to determine if that's appropriate for them.
- We'll give information and guidance and education to go along with that. There was a fact sheet published recently by CISA, NSA and NIST that had a lot of steps for things you could do to prepare while the standards are just about ready.

The NCCoE Migration to PQC Project

- The NCCoE project is focused on the migration and have partnered with over 30 different industry partners. They're developing tools that will help with migration, hosting webinars, and publishing reports.

Conclusion

- The standards are almost here, which will be a great milestone. We are grateful for all the help from the community. We have a webpage, and you can always contact us. This will have a lot of announcements and information.

Discussion

- **Mr. Lipner -** Just a comment, the diversity of algorithms that you're standardizing feels like good news, bad news. The bad news is that it's going to leave people with some uncertainties about what to use and how to proceed. The good news is that you're probably going to force this issue and that's a long-term benefit.
  - **Mr. Moody -** We tried to signal that Kyber and Dilithium are the main two that we expect them to use.
- **Mr. Lipner -** The points we were discussing before, people realize that they may have to adopt one or the other. That will encourage them to make the investment.
- **Mr. Venables -** The questions are getting a bit too much about the future.
  - The urgency is a function of:
    - When we think a cryptographically relevant quantum computer will be available as a function of two things, the engineers building such a thing and the number of logical qubits that are needed.
    - But it's also functioning in algorithmic campaign space.
      - In the past year, there's been several papers, particularly coming out from China, that made amazing claims about how they've reduced the number of qubits needed.
      - Most of those appear to be debunked quickly, but I would expect we'll see more of that.
      - We have our own people (to look at these papers) for large companies, but there's probably a lot of nations that don't have cryptographic expertise.
  - Should we be looking to NIST to look at those papers and put out advice, or is that a DHS thing or NSA thing?
  - **Mr. Moody -** It's not NIST so much that provides public comments on those algorithms, other than informal networks where we communicate with each other.
  - **Mr. Scholl -** I don't know. We generally do not do a public peer review of papers and then provide our opinion. The community is robust about having analysis and discussions when those multiple papers have come in.
    - On the quantum side, the Quantum Economic Development Consortia (QEDC) is an external organization that is very good in helping us to understand the quantum engineering side. Internally, we have our own physics laboratory that we often lean on.
- **Ms. Flynn Goodwin -** Question, at what point will this issue hit the tipping point where we must translate it for real people? Standards are years away. At some point, this is going to be a real cyber hygiene issue for small and medium sized businesses who are still running on-prem software updated twice from the time they bought it.
  - Are you going to be the ones that explain all of this to people about why they should care? When do we start that conversation?
  - **Mr. Presman -** Good news is that some of this should be relatively seamless to the public. When Google or Microsoft updates its algorithms, to the regular population that's going to feel just like every other patch or update that's ever happened.
    - In the general population, most people aren't running legacy and hybrid systems.

        - What's going to be more of a challenge is in critical infrastructure, where their OT may have units that are 30 years old and don't get updated on a regular basis. There's not a cadence like that in the OT world.
        - CISA has the lead on working with critical infrastructure and has already done many engagements with each of the 16 critical infrastructure sectors. It's not going to be enough because we know that the critical sectors are quite fragmented. But at least at the leadership level, we know that those critical infrastructure sectors are being engaged and beginning to think about the implications.
    - **Ms. Flynn Goodwin -** It's about the overall scheme or risk priority. When CISA and NIST have to think about what risks we are going to care about, where does this sit in the stack? We talked two meetings ago about the technical debt that the United States must pay down as we migrate from legacy things to new things. This is just going to add to that technical debt burn down that companies will have to spend as they're migrating.
        - Where are we going to put this in that stack frame, and how are we thinking about that as a nation as we're encouraging companies in their risk management profiles to go and make priority decisions? That seems like a conversation, and a priority setting to be pushing now.
    - **Mr. Lipner -** Two other things to reemphasize. It's not only legacy systems in enterprise customers. And you're right, they'll get a product and they'll be done. There are lots of un-upgraded small businesses or systems running Windows XP, on SP2, if you're lucky.
        - Maybe what you want to do is to get the Small Business Administration to undertake this as a project. Not sure whether that's something that would be in their remit, but it's part of the economy and it can have a negative economic impact.
- **Mr. Gantman –** A quick question or comment. Considering that our remit here is non-classified federal networks, how much of a concern is harvest and decrypt for non-classified data?
    - **Mr. Presman –** There is a significant amount of unclassified data that's sensitive and that will still be sensitive in 15 years:
        - Personal health data, law enforcement data, intellectual property.
        - One thing to think about in terms of quantum computers is these are large, very complicated instruments, not likely to be in the hands of mom-and-pop hackers. More likely to be in the hands of the national threat actors. When we think about what their interests are, that puts that more in the crosshairs of what are the concerns on the unclassified side. Intellectual property is one area where there is a lot of interest for those persistent actors. That's usually unclassified information but highly sensitive.
- **Mr. Lipner -** Basically forging signatures, can you break the key that somebody signed with?
    - **Mr. Venables -** Store now and repudiate later:
        - There is a significant amount of sensitive information in the private sector that should be worthy of consideration for decrypt later, but the larger concern in the private sector would be the integrity angle, not the confidentiality, which is store now and repudiate later.
        - Imagine somebody digitally signing a contract today that has a 30-year binding, and this could be repudiated 20 years later, and it still has a lifetime.
        - We may be underestimating the store now, repudiate later.
- **Ms. Fanti –** What about the situation with hardware, in particular hardware with firmware, that cannot be updated. Do you have thoughts on that? Estimates on what that cost would be for updating that hardware? Replacing it?
    - **Mr. Presman –** We do not have estimates on a national level. Within the federal government, downline estimates that exist around the cost of paying down this tech debt that we've been talking about and upgrading legacy to more modern systems. It's not cheap.

  - **Ms. Fanti -** What does the schedule look like? At what point do those devices get updated? Should everything that they interact with also be updated at that point, or maybe you build in some kind of versatility?
  - **Mr. Presman -** There isn't an answer. It's a good and important question that needs to be constantly brought up, but it's not one that has an answer at this point.
- **Mr. Lipner -** Just one other thing that is worth thinking about. We did the National Academy study on the future of encryption, one of the things that came up was improvements in classical factoring, where, even if you don't have a quantum computer, somebody could, at some point, make a mathematical improvement that would threaten RSA.

Cybersecurity for USG Research and Development

Terry Carpenter, Deputy Office Head, Office of Business Information Technology & Deputy CIO, NSF
Linnea Avallone, Chief Officer for Research Facilities, NSF

Introductions

- **Mr. Carpenter –** Thanked the board for the opportunity to speak on behalf of NSF and the mission they have. Linnea Avallone runs the facility site, which is an important part of the questions you're asking us.

Question: As you're looking at the NSF approach to the ecosystem of research, and how do you think about cybersecurity across that ecosystem?

- There's been a couple of studies that have shaped this philosophy.
- We generally avoid being prescriptive because it's the responsibility of the awardee, by the way in which we issue funds, to support research activities.
- We do have oversight accountability and we're constantly trying to improve the way in which we think about performing that function of oversight and making sure that we assure cybersecurity and resiliency through the research ecosystem.
- Some principles that we think about are that cybersecurity is changing every day. We must do this in a way that adapts to the latest changes in policy practices, and our adversaries have a vote.
- Three buckets or areas of interest:
  - The pure research domain: what we're asking people to think about cybersecurity,
  - Cooperative agreements domain: We're asking people to conduct research in facilities as a group or consortium,
  - The agency: what it does as a foundation and how we keep our data.
- This programmatic approach is one we're really thinking about, how does that change now as the lines between those three areas get blurry? These are huge national assets.
    - Facilities that are out there performing and enabling this kind of research,
    - The telescopes and other national assets that we have,
    - We need to make sure that everything we do helps them to be secure.
  - We don't operate them. That requires a significant level of engagement, especially with these larger facilities where there's a structure, a team, behind the operations of the facility itself, engaging with the leadership, ensuring that the processes, the best practices are all implemented.
- And finally, as we think about awardees, we think about that timely information. Things are happening so fast, getting them to provide us timely information is just as important as us setting the stage early in the process, so we know when things change.

Research Facility Resilience

- With large research facilities, research facility resilience is key. The sophistication of today's activities against us are becoming more planned out, well executed, and it's getting harder to ensure that level of protection.
- Everything we're doing in the facility is to minimize the likelihood, the impact, and the disruption, and ensure the integrity of the science being performed.
  - Disruption can be a hard one, when you think about multiple tenants inside one of these large facilities. You invite people in and there's only so much time. It's almost back to the days of computer scientists time sharing on the large mainframes. We invite them. When they get their time slot and perform their research a lot of things must come together.
    - Any disruption in that schedule could be very impactful to the outcomes that we're all looking for on a national level for the research work being performed in the facilities.
  - Looking at it from a CIO perspective, using zero trust architecture as a risk management framework, things that have been defined, proven that we are constantly implementing and changing the way in which we're implementing and monitoring these frameworks and standards that help us do a better job. It's important as we think about this research ecosystem, that we apply these things.

What are the requirements for NSF research programs?

- When we think about programs with facilities, specifically the major facilities, that are not managed or owned by NSF, these are operated under cooperative agreements. There's a legal binding agreement in place that places the expectations of cybersecurity standards and policies against the operator. We are constantly updating our expectations, our requirements, our review process of how we think about helping those operators to assure cybersecurity.

- **Mr. Groman -** Asked about the awardees.
  - **Ms. Avallone -** It's a mix. When we're talking about our major facilities, in some cases, it's a university or a consortium of universities. We work with quite a number of university consortia and in some cases, we work with not for profit corporations as well.

- When we think about the facility, we look at what we require them to do.
  - One of the first things is in the proposal process.
    - The facilities, when they get an award, are required to provide an initial cybersecurity plan within 90 days. We also have an element that is part of the formal review process and there are some annual reviews that happen and some requirements against those where we're looking at the performance of cybersecurity throughout the effort that's ongoing in that research,
  - In the NSF Proposal and Award Policies and Procedures guide, Section 2 talks about the proposal process and what's required.
    - We require a data management plan.
    - The idea is to understand how they are thinking about and protecting that data.
  - Section 9 reflects the latest of the NSPM 33 guidance (https://www.whitehouse.gov/wp-content/uploads/2022/01/010422-NSPM-33-Implementation-Guidance.pdf), which is making sure that the national security information that we need is also included.
    - It's focused on the disclosure of support and in-kind contributions.

    - What that really means is what are they getting and what are they using outside of themselves to perform this work? Are there key personnel involved in supporting the research work, and how are they disclosed?
  - If they didn't disclose it in the proposal, these two requirements require them to update us immediately or within 60 days if they have a change in that in-kind support.
    - There are provisions there to help us try to improve that information flow to know when things are different than what they proposed.

- **Mr. Groman -** Asked if these are awards or grants under federal law?
- **Ms. Avallone -** It's a combination. Some background contexts:
  - NSF is somewhat different from other federal agencies. If you look at our budget, which is around $10 billion, roughly 92% of that money goes out the door to support research, whether to research facilities, to individual grants, or to fellowships. In a typical year, our facilities operations budget is about 1.2 billion, but our grants are about 6 billion.
  - Many of the awards that we make in a year go to individual faculty or researchers, usually at universities, but also, at museums, not for profits, and other types of organizations.
  - The research facilities we typically run by cooperative agreement, which is another type of financial assistance. We have a couple that we run by contract, but primarily through cooperative agreements.
- **Mr. Groman -** So, requirements that apply to a federal contract where a researcher under a contract must comply with FISMA and RMF does not apply to grantees.
  - **Ms. Avallone -** That's correct. These are not federal institutions. We do ask for due diligence. Awards formally do not go to individuals; they go to institutions. It's the institution that is on the hook to comply with our grant terms and conditions.

What are the guidance that policy programs that a research agency might be able to leverage from NSF?

- We went out and looked at some of the programs. There's a lot going on.
- There are three large programs and investment areas that we have in cybersecurity:
  - CI compass, which is focused on cyber infrastructure and data management research.
  - Trusted CI, which is a consortium focused on cybersecurity and cybersecurity assessments. That group just had a conference out on the West Coast this week, looking at bringing people together to figure out how to help those institutions do a better job of raising the bar across all of them together, rather than having each one must figure it out themselves.
  - Regulated research, which is a network of about 270 institutions, helping each other to implement affordable and effective cybersecurity compliance programs at academic institutions.
- We are supporting these efforts, funding some of these activities that then help them to elevate their game at the institutional level.

- **Mr. Groman** Regarding incentives. His recollection is there was not a robust accountability mechanism. For example,
  - The university received a huge amount of federal data from NIH or elsewhere, and there was a breach, and all this data is compromised. The university had no liability or responsibility at all. The university said, "I'm not taking responsibility. If you want me to do that, then I can't do research."

  - They didn't have any incentive to bolster security, because if something went wrong, they had no accountability or liability. Does that make sense as a question? What happens if they experience a breach?
  - **Mr. Carpenter -** I don't have the answer. We would need to consult with our general counsel to look and try to dig into that.
    - It's a good point. I don't think about the incentives of a research institution, but there's a lot that goes with not getting a bad grade. We're constantly putting resources and funds out there to academic institutions to perform this work.
    - The proposal review process would be the only lever I see right now.
  - **Ms. Avallone -** We rely on 2CFR 200, which is the uniform guidance for most of our award conditions. There may well be language in there that gets to the question that you're asking.
  - **Mr. Groman -** And then do we do due diligence on these entities? Do we follow up?
  - **Ms. Avallone -** So since I work on the facility side, I don't have direct knowledge of what folks are doing on the grant side. There is an annual review process for every award that the agency makes, so there are those opportunities for follow-up. As for certifications: there are a number of certifications that we ask for when a proposal is submitted. I think there are potential leavers, I just couldn't say exactly what we use in that situation.
  - **Mr. Carpenter -** (to the best of his knowledge) There's no entity inside of NSF that goes out there and does that. Not like the Department of Defense where there are different levels of things. Again, we're talking about unclassified efforts here.
  - **Mr. Groman -** We just made this point of an unclassified effort still consisting of data that would be highly valuable to an adversary. For example, genome sequencing of 100 million Americans would produce something medically classified and valuable. This merits a deeper look.
  - **Mr. Carpenter -** We're starting to ask ourselves questions like, if you look at one piece of a set of data like that, you can assess its risk easily. The risk changes when there’s an aggregation of the data. This is a good point.

- Trusted CI group website: https://www.trustedci.org/. There is a lot of good information on their site.
- The NSF program: Funded in our computer and information sciences engineering group. It's a cybersecurity innovation for infrastructure (https://new.nsf.gov/funding/opportunities/cybersecurity-innovation-cyberinfrastructure-cici).
  - It's study and research work to learn what might be new out there. Also to help people plan on how to move into that new, better stuff.
- We're thinking about the work we're doing in multiple layers, not only what's the next big thing, but how do you help people to get into it? How can small academic institutions be able to compete and be just as secure.
- Another investment area, Advanced Cyberinfrastructure: (https://www.nsf.gov/div/index.jsp?div=OAC)
  - Again, this is more in that research realm of advanced cyber infrastructure investment. Examples of programs out there that are looking at specific topics within this space of research, infrastructure, security.

What are some issues that drive NSF?

- Stewardship for national assets is critical and drove a lot of our thinking.

    - If you can't get that exact moment in time when you need to be able to look up in the sky and perform that research, when's the next time that it comes around? And what did we miss? What opportunity cost did we lose there in having that happen?
- The growing concerns over national security at federally funded research, unclassified data sensitivity, and NSPM-33 (https://www.whitehouse.gov/wp-content/uploads/2022/01/010422-NSPM-33-Implementation-Guidance.pdf).
    - The idea that this work is being done in an independent way, but as part of the heterogeneous nature of major facilities coming together to provide that capability for others.
    - It resulted in a couple of things.
        - Commissioned a MITRE JASON Project study in 2020. It was completed in 2021.
            - A lot of the things that you're seeing in this brief stemmed from or were enhanced by the learning that came out of the JASON study.
            - We also established an office that reports to the director on the chief of research security strategy and policy.
        - Open research environment
    - Cultural differences between higher education and the industrial base.
    - Came from DoD and when we think about the culture, you have a core mission difference and a profit for the defense industrial base. You have peer review, intellectual property.
        - Everything from research environments for integrated teaching to restricted staff, decentralized, project based, collaborative, centralized span of control.
    - This idea of open research has a big impact on being able to advance some of these hard topics. You need to have an open research climate where we can share and learn from each other to help accelerate the outcomes that we're hoping to achieve.
- **Ms. Flynn Goodwin -** In our community, there's a lot of research that's done, not in the traditional ways that you work, but by independent researchers. How do you think about that culture, the learnings that get reflected from that culture, and how that reflects into the more traditional approaches of how you think about learning from an NSF and cybersecurity perspective?
    - **Mr. Carpenter -** NSF is supporting the ecosystem that NSF directly funds. In turn, a lot of these activities that we're funding, like the centers and networks that are working towards this information, has applicability other places, and can be used other places. We don't know how or if it is being used in other places. There's a long chain there and there may not be an answer. Our hope is we're putting the monies in the right place that not only affects the stuff that we directly need to manage but improves the overall ecosystem of science.
    - **Ms. Avallone -** Many of our staff who manage the programs are practitioners in these fields.
        - One way that we do what you are asking is on an individual-to-individual basis. The program staff go out to conferences, they know their discipline very well, and networking happens at that level rather than at a formal agency level.
        - We rely on the disciplinary expertise and networks of our staff to keep us informed about what the trends are, where we should be thinking about developing new programs, and taking in research results, not just from what we fund, but from the broader swath of the community.
    - **Ms. Flynn Goodwin -**This is a space that prides itself on counterculture. Thinking about a way for the NSF to solicit opinions directly from that counterculture might be an interesting way to get fresh takes on well-established challenges.

- **Ms. Avallone -** So I can say that every one of our directorates, including the Computer and Information Science and Engineering Directorate and the sub office called the Office of Advanced Cyber Infrastructure, have advisory committees like this one.
  - Don't know the current representation on those committees, but that would also be a way to get some of the different cultures involved.

What is NSF doing about new guidance?

- This is a rapidly changing area for us because of some of the results that we learned in the JASON (advisory group) study.
- The Research Infrastructure Guide for Major Research Facilities will be updated in 2024.
  - There's going to be a new section on cybersecurity in it.
  - There is a draft out there for comment.
  - It recognizes that this area must adapt to new challenges, support the research workflow, and build out this resilience where we're not impeding science because of something that happened, and not impeding on national assets.
- We stood up that new office a couple of years ago and we established a new position as a cybersecurity adviser for research infrastructure to help advise and shape guidance.
  - This is an expert in the field that also came from the academic culture and understands how they're implementing cybersecurity in the academic community.
  - Three responsibilities for the new position
    - To ensure completion of the Jason report recommendations,
    - Help us refine the posture of those major facilities, and
    - To resource for program officers and research facility operators, those that are doing the work, and making sure they're secure.
  - **Ms. Avallone -** This position reports to me. The focus is initially on our major research facilities. As it evolves, the person will also work with our chief of research security strategy and policy.

Discussion

- **Mr. Groman –** Asked if there is a security and privacy training requirement for anyone who works on one of these grants?
  - **Ms. Avallone -** At the grant level, there are some basic requirements.
    - Ethical and responsible conduct of research.
      - That's a requirement of every institution that receives an NSF award.
      - Does usually touch on privacy issues.
      - Probably light on formal cybersecurity.
    - In general, our major facilities have several requirements for the types of personnel that need to work on these projects, and we have formal expectations for their core competencies, including in cyber infrastructure and cybersecurity.
  - **Mr. Groman -** Your observations that our threat actors have changed exponentially in the last five years are correct.
    - It’s not if but when your databases will be hacked.
    - More consistency is warranted given this ever increasing and rather dramatic threat.
    - We want the research to go forward, but sometimes we say “no” to researchers: you can't always extract your data set, you can’t always take it home, you can’t always work on your own device.

- **Mr. Carpenter -** One of the things that the CHIPS and the Science Act helped the foundation to do was to stand up the first new directorate in 30 years called the TIP, Technology, Innovation, and Partnerships (https://new.nsf.gov/tip/latest).
    - What's interesting is it's helping us see and learn how we can, in accordance with the laws and policies, get funds out there for research to be conducted in a different way.
        - Not economic development.
        - Looks more like relationships in a tighter ecosystem.
        - Funding things we call engines and accelerators. The idea is to get a little more influence over how that work is going to be conducted. These accelerators, all part of the ecosystem for that topic area, will come together, helping to advance science in a larger picture, but also in a way which we could influence some of these controls,
    - Looking at how to make sure these partners who are participating in this slightly different model may be able to inherit some things that are at the engine.
    - May help small partners if they come and participate.
- **Ms. Avallone -** One of the challenges is that we have limited authority when we make grants; there are only certain things that we can ask institutions to do.
    - We have certain requirements that we ask. Often, we ask the institution to self-certify that they're doing certain things, rather than prescribing that they do them and how they do them. To delve more deeply into this, we would need one of our attorneys who specializes in the grant management to talk about what we can and can't do.

- **Ms. Moussouris –** May be out of scope from what you're presenting but universities are subject to export control. How are you integrating with the enforcement of export control in the context of technology advancement integrity?
    - **Ms. Avallone -** That is a responsibility on the awardee. We do not engage on that at all.

The Chair recessed the meeting for a 60-minute lunch break.

NIST Cybersecurity for Research Guidance

Connie LaSalle, Senior Technology Policy Advisor, NIST

Introductions

- **Mr. Lipner** Introduced Connie LaSalle from NIST
- **Ms. LaSalle:** Sits within the Information Technology Lab (ITL). Background is largely in policy and federal operations. Has worked in the private sector and has been at NIST for the last two and a half years.

Background: §10229

- Will talk about one effort that is new for us as of the designing of the CHIPS and Science Act[1] in August of 2022.
    - §10229 is the specific area with the details of the mandate. This is not cybersecurity research, but this is cybersecurity for research. §10229 is within the CHIPS and Science Act.
    - CHIPS part: bringing domestic manufacturing back to the United States, specifically semiconductors.

[1] FACT SHEET: CHIPS and Science Act Will Lower Costs, Create Jobs, Strengthen Supply Chains, and Counter China https://www.whitehouse.gov/briefing-room/statements-releases/2022/08/09/fact-sheet-chips-and-science-act-will-lower-costs-create-jobs-strengthen-supply-chains-and-counter-china/.

  - Science part: included authorization language and requirements on dissemination of cybersecurity resources for research institutions, in particular research institutions within higher education receiving $50 million a year or more in research funding.
  - NIST has been asked to investigate resources that can help them identify, assess, manage, and reduce cybersecurity risks related to conducting research.
  - This is important because universities are small cities:
    - If they have a clinic on campus, that is PHI (protected health information) that they must manage.
    - If they are offering a student aid in some way or as their financial service provider, etc.
    - There are a lot of different regulations and ways that you can interpret the kinds of information that a campus could be responsible for.
    - Because research is highly decentralized and there is not one pot of cybersecurity funding that gets dispersed down into the research projects and programs, it can make for a complicated ecosystem.

§10229 Implementation Timeline

- Starting in January of this year, we started by getting to know the topic. We started by getting to know the professionals in academia who are already working to solve these problems.
    - We then issued a request for comments on the state of cybersecurity across your research portfolio today.
    - We engaged one on one with organizations. We spoke with several associations in DC, that are intimately familiar with the regulated research ecosystem.
    - We talked with other cybersecurity companies and service providers that specifically focused on higher education.
    - We then decided to put together a NIST Interagency Report (NISTIR) to summarize and document some of the challenges, risks and approaches that people are taking to manage those risks across higher education.
  - We are now in a comment period that ends October 31.
  - Much like other cybersecurity challenges, it comes down to budget, workforce, and where can this play a particular role? Whether just highlighting the resources that we already have or creating new ones?
- **Mr. Lipner –** Asked for the link for people who have not seen it.

Community Engagement To-Date in 2023

- We have had direct engagements with 20 individual institutions. Had brokered public engagements through EDUCAUSE, an association, that is highly engaged on these topics. Also had a couple of group discussions in addition to one on ones.

Community Feedback Analysis and Key Questions Informing NISTIR 8481

- Our approach has been to take all his feedback, put it in this one document where you can see everything laid out nicely. The contents that you can expect responds to five core questions.
  - First, what are the common and cross-cutting cybersecurity challenges and risks that higher education institutions are facing when it comes to their research?
  - Second, are there unique challenges and risks. We know that the Hubble telescope is going to be managed differently from an iOS device. How do you look at unique fields of study and provide each of those audiences with practical cybersecurity guidance?

  - Third, what resources already exist that you are using, whether they are NIST resources or something that your campus CIO has put together?
  - Fourth, are the resources sufficient? Could they be better, or could they be replaced with something else?
  - Fifth, if NIST were to play some kind of role in developing new or tailored resources, what would that look like? Who should be involved? Is NIST leading the pack? Are we simply one member of it? Is there somebody in higher education who is really leading, and this role can be to highlight the good work that they are already doing? There is a role that we can and should play, help us figure out what it is relative to your community.

Summary of Feedback: Common Challenges & Risks

- Six big categories on how people responded to that first question of common and cross-cutting cybersecurity challenges and risks.
  - Workforce. The workforce issues are hard to compete in general. It is particularly difficult for smaller institutions.
  - General awareness and cyber hygiene are something that they should and can do better.
  - There is also a culture clash. Whereas centralized enterprise organizations believe that they can adopt an approach and enforce it down, that does not work in academia. This is a highly distributed portfolio when you are looking at research.
  - Limited budgets. It is largely dependent on the sponsoring organization or funding of these requirements. The budget that is centralized for campus Wi Fi, for example, is not necessarily the bucket that will be available for things like the devices need for the research. There is an open question of how some of the approaches do that work for the education mission.
  - It is also a complicated requirements landscape, particularly when you have two different funding organizations interpreting a requirement in different ways. Whether that is what CUI is or is not or the scenarios in which certain NIST SP 800-53 controls apply, it is very heterogeneous.
  - The rapid pace of innovation. If a bank such as JPMorgan Chase is struggling to figure out its approach to AI enabled threats, you can imagine that a small school in Arkansas is also struggling with that.

Summary of Feedback: Unique Challenges & Risks

- In response to the second question, these fields of study where unique cybersecurity risks come into play.
  - The needs for infrastructure, devices and the technology look different and the thinking about the drivers of this kind of initiative.
  - There are different threats and different entities that are interested in the kinds of outputs of research associated with these areas.
    - Engineering, quantum computing biotechnology, etc.
    - These are things that have a great potential to increase GDP in a country if you have a commercialization engine that leverages that research and enabling competition at both the national level and global level.
  - It was very interesting that these were some of the specific fields of study that were noted as having unique considerations.
  - The hope is that through this current public comment period, we will hear of more.

Summary of Feedback: Future Work Suggestions

- Academia provided suggestions for future work that NIST could take on; general suggestions that somebody could take on in the future. One example noted across respondents was this idea of more shared infrastructure and shared services.
- An opportunity and a challenge is who should be responsible for managing a secure enclave?
    - Does it make sense for the federal agency that is funding a program and is equally set up to benefit from the cybersecurity requirements and privacy requirements that are implemented?
    - Should they take on more of a responsibility by providing the infrastructure that then can be shared across performers so that they are not worrying about trying to get to the exact same level of due diligence across the projects that they are funding?
    - Similarly, is there an opportunity for separate stakeholders, such as the funding agency and the higher education community, to come together to establish some sort of shared services or shared infrastructure?
- One big theme targeted cybersecurity resources and trainings for the research community. Often cybersecurity awareness training is built for a corporate or federal audience, which does not always translate for researchers.
    - Open question: can you talk to a research community instead? And how would you do that? Who would do that?
- Collaborative engagement with other federal agencies that are imposing requirements and interpreting guidance, and engagement between funding organizations and the implementing organizations and academia. There was a request for an open line of communication and continued engagement on this topic.
- NIST has asked for guidance about how to apply existing guidance.

- **Ms. Moussouris -** What about the sort of shared services that are more centralized?
    - **Ms. LaSalle -** I asked that same question of the folks who are setting up a secure enclave for their campus. I think they noted it, but their point to me was this versus nothing at all. We are going to try this out, see how it works, and then note that we might be creating a single point of failure.
    - **Ms. Moussouris -** It is just that not so long ago, when I was working, defending an academic system. We were getting attacked by China and did not see that was happening. It was easier knowing that they had to go machine by machine, as opposed to having all the research in one place. We had a chance to not have them get everything.
- **Ms. LaSalle –** Great point
- **Ms. Fanti -** Have you heard about the Cloud Bank program through NSF? If not, it might be worth looking into. It is a program where academics, who are funded through certain calls, can get Cloud Credits. And they can spend those Cloud Credits on different major cloud service providers.
    - **Ms. LaSalle -** I do. I am aware.
    - **Ms. Fanti -** We need to know how that impacts cybersecurity for computing related research.
- **Mr. Lipner –** Regarding the list of 6 concerns: My impression is that the big one to overcome is the culture clash. Did anybody engage with that?
    - **Ms. LaSalle -** I think that is a practical reality as far as the incentives for getting researchers to adopt any kind of practice.
        - There are some universities who are standing up internal lines of business, trying to make the case that, by going through their trusted inventory of applications or by using their internal consulting service for their program that makes it their value added.
        - Some organizations recognize that there is a power imbalance, that they cannot just set the university-wide policy and trust that the researchers will follow it.

        - Example: SecureMyResearch program at Indiana University.
            - They pitch themselves as a value add resource to improve your research.
            - To make sure that you get the right of first authorship.
            - Provides some level of provenance for any findings, it is helpful for supply chain purposes. That can also be helpful for voting,
        - Part of this too, is meeting researchers where they are culturally versus trying to force them into a top-down command and control model.
        - Also looking at the organization chart is telling, where you have the Vice President of Research who might be taking more of a compliance view versus more of a risk-based approach. That is where you get the same old culture clash, whether it is the commercial sector or Government.
    - **Ms. Moussouris** - I just want to make a comment that is interesting and innovative, tying the "you get first authorship or primary authorship," as a value add. I think that is aligning the motivations of researchers with the goals.

Potential Next Steps for NIST

- We have three buckets of things that we could do, and we are actively asking for feedback.
    - Community specific cybersecurity resources,
        - This is something that is probably going to be welcomed. It is a matter of getting specific about what that could look like, what the scope would be, what community we should engage with, is it based on that unique field of study? For example, do we need one profile for quantum scientists everywhere or do we take more of a generic approach to start in the hopes that we have some practices that, no matter who you are, you can at least be thinking of coordination.
        - That is a simple request to keep the lines of communication open between academia and the federal government, but it is also requested the federal government would coordinate more.
    - Have been in coordination with the NSF team from the last session. We both talked with OSTP (Office of Science and Technology Policy), which is leading the charge on NSPM 33[2] implementation.
        - This is generally touching on research security for the entire federal research and development enterprise, not just scoped to academia.
        - Ultimately, that guidance will point to the NIST cybersecurity program.
        - There will be this nice circuitous policy situation where, whatever we figure out from a bottom up, federal to academia, engagement will be doubled down into OSTP guidance.
        - If you have any questions about that, I know it's not the scope of my chat today, but we can talk about that.
    - Workforce and budget constraints.
        - The potential is to look at what NIST does through our NICE (National Initiative for Cybersecurity Education) program, and through other efforts that could be used to help with capacity building.
        - We don't see ourselves as the ones operationalizing the training but can help advise on the curriculum.

[2] GUIDANCE FOR IMPLEMENTING NATIONAL SECURITY PRESIDENTIAL MEMORANDUM 33 (NSPM-33) ON NATIONAL SECURITY STRATEGY FOR UNITED STATES GOVERNMENT-SUPPORTED RESEARCH AND DEVELOPMENT chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/https://www.whitehouse.gov/wp-content/uploads/2022/01/010422-NSPM-33-Implementation-Guidance.pdf.

Appendix A. Existing Cybersecurity Resources

- A big piece of this requirement was to disseminate resources.
  - We have created this appendix showing what already exists and increase awareness of what is available to them.
  - You can see things like the:
    - CSF (Cybersecurity Framework)
    - The Privacy Framework.
    - The Risk Management Framework and Special Publication (SP) 800-171.
  - Beyond just the documentary guidelines, we also point out that we have the National Cybersecurity Center of Excellence (NCCoE), which could be a hub to bring together disparate research communities to have a conversation about cybersecurity.

Areas Where We'd Like Feedback

- The last slide shows areas that we highlighted in our note to reviewers for the NISTIR,
  - Can you validate that the cybersecurity challenges and risks that we have captured track for you? Do we miss anything? Did we miss some nuance?
  - Similarly, as far as the potential next steps are concerned, NIST is also budget constrained. If we are going to dedicate effort, where are you going to see ROI for every hour of effort from this?
  - Finally, do you have a cybersecurity resource we should highlight? Or do you have some commentary about a resource that we are now promulgating through Appendix A. Let us know.

Discussions

- **Ms. Fanti -** There have been a few incidents where researchers do not engage in best practices and minor incidents occur as a result. Also concerned that having these researchers go through additional cybersecurity training is not going to have any impact.
  - One thing that might be useful instead, maybe you have a short checklist of the bare minimum of things that everyone should be doing, like enabling two factor authentication on all your email accounts or using a password manager. It is something short and to the point without having to spend a lot of time.
  - **Mr. Lipner -** The last thing to inflict on universities is the RMF or a long list of individual compliance things. But in terms of affecting what people do…
  - **Ms. LaSalle –** One idea was that offering something like this, but by data type, could be a nice dichotomous key for researchers. If they have been told, you are handling this CUI, therefore, double check that you have these practices in place. Does that sound like something that would work?
  - **Mr. Groman -** Doesn't want to remove the requirement that someone think and understand the risks of their project. However, what you are doing is a short form of a baseline risk assessment for them.
    - You might have something like, if there are other factors or criteria beyond this, these might be other things that you consider.
    - It will certainly be helpful.
    - Then the upfront arguments over what data is entered is what will have to be had.
    - Something useful is to try and distinguish datasets. Not just if you have SSN but if you have 100 million of them or have 1,000 of them will also distinguish the kinds of security you might want to have.

        - Or if you have profiles on a million people, but each profile has three data points, or each profile has 100 data points, the risk of that data set is dramatically different. But you could certainly do the beginning of helping them evaluate the sensitivity of the data set, and therefore, that would implicate controls. Does that make sense to anyone?
- **Mr. Lipner -** Another thing, if you can co-opt, senior research officials from universities and get the group of folks who will speak about these issues within the committee as insiders. That might be helpful, particularly with the culture issue. The problem that you, and the people who care about security within these institutions, are over trying to overcome, is "I cannot be bothered with it. I have research to do" or "I can't be bothered with that. I have research to do, and I have tenure."
    - **Ms. LaSalle -** That is exactly right.
    - **Mr. Groman -** At some point, no one wants to be having these discussions because we want to restrict the ability to do extractions of massive datasets. I was going to slow down research and millions would die, because of the security controls. At some point decisions must be made.
    - **Ms. LaSalle -** Do you see privacy enhancing technologies as a path to meet that kind of balance the need for research and minimization?
    - **Mr. Groman -** First, this is not about risk elimination, because that means data augmentation. It is all about risk mitigation.
        - What we want, for those projects that are sensitive because of the volume and quality of data or potential uses, is to have controls in place that will mitigate that risk and have them implement it.
        - Some of the easiest ones are on credential, MFA, password controls, and do not let people in kindergarten get access, and things like that can be done.
        - It will always run into cost and convenience.
        - When I was the CPO of the FTC, we had an economist that did antitrust studies on hospitals. We would get datasets in production of patient records on 50 million people in the hospital system. There was a time when that would go into their home on an unencrypted laptop that they owned. At some point we had to say "no."
    - **Ms. Moussouris** – Circling back to tying it to research and motivations, I remember the controversy around who discovered the HIV virus. Keeping the data private and secure and tied to specific identities, it would allow, not just from the enclave of researchers working on that in a localized university sense, but a country-to-country sense. As I recall from that controversy, it was "we think we found it." "So do we." "Send us your sample and we can confirm it". If this had been in place, in this case it was the US researchers that were the sneaky ones. There is precedent for this kind of application.
    - **Mr. Groman –** When you do research at this level, you must be able to establish the integrity of your research. Once the data goes in, you cannot delete it, you cannot alter it, it cannot be removed, because you need to be able to reproduce the results of your research. That goes to data governance.
        - That piece that we need to add in addition to having logs, there should also be something that stops it from going in and out when it is not authorized. It is that extra layer that we need.
- **Ms. LaSalle –** That resonates so much. During a regulated research community of practice meeting out of the CIA triad confidentiality, integrity, availability, I asked what are they most concerned about?
    - They are focused on confidentiality because they were told that they must be. I expected integrity, or maybe availability.
    - **Mr. Groman -** If it is an FDA project, integrity is critical. For example, you want approval for a new drug. Integrity is critical or you are not getting approval. You must be able to reproduce it,

depending on the kind of study. All studies are not the same. We do not want a one-size-fits-all that everyone must do because these do vary. The problem is that when you do not have a one-size fits all, that requires someone to conduct the analysis and be thoughtful, and then own it.

- **Ms. LaSalle -** You and I are on the same page.
- **Mr. Groman –** The ownership for the accountability piece.
- **Ms. LaSalle -** The answer I got of "confidentiality" demonstrates the culture clash and the mismatches in motivation that is happening. The extent that we can help cybersecurity workforces in higher education provide better services to the researchers in a way that focuses on those motivations, to include integrity, because that probably should have been the number one answer for many cases. That that might be an element of the direction that we go.
- **Ms. Moussouris -** As a molecular biologist, the confidentiality bit I worked in an AIDS and hepatitis research retrovirology lab and, from PhD to PhD, they did not want to share. The only reason I knew that they were working on similar things or that one was trying something that the other had already failed, was because I was an undergrad and therefore could not publish and so I was perfectly safe.
  - Surely that bit about confidentiality being the primary. Yes.

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

CSF 2.0 Update

Cheri Pascoe, Director, NCCoE, NIST

Introductions

- **Mr. Lipner -** The speaker for this session is Cheri Pascoe, and she is the new Director of the National Cybersecurity Center of Excellence (NCCoE). Congratulations, and thank you for joining us to talk about the cybersecurity framework.
- **Ms. Pascoe -** Thanks so much for having me.
  - Was the prior lead for the NIST cybersecurity framework program and is also the architect for the CSF 2.0 draft.

Overview

- The CSF is widely used around the world, by organizations of all different sectors, different sizes, and now it's become widely viewed as essential and foundational to reducing cybersecurity risks.
- It is a big deal for us to proceed with an update. We want it to remain relevant to address the evolving cybersecurity challenges, keeping pace with the standards and technology landscape.
- We want to make sure that we are not losing anybody that is currently using it, and that our audience is growing.

Governmental Policies on CSF

- The CSF has become known as kind of a voluntary framework, which is partially true.
  - Federal agencies are required to use it, but increasingly, we are seeing mandates around the CSF, especially overseas.
  - In response to the NIS directive, Italy and Poland have adopted the CSF as a requirement for use by critical services,
  - In a recent cybersecurity strategy that came out earlier this year, there was quite a bit of discussion about cybersecurity regulation harmonization. As part of that it talked about how

federal regulators in the United States can use the NIST cybersecurity framework, as well as the CISA CPGs, to advance that regulatory harmonization.

- NIST has been serving as a technical adviser to the FCC, which runs the forum of all the different federal cybersecurity regulators. We meet weekly to carry out the direction of how regulators can use the CSF to inform their regulations and to explore the existing gaps in the regulations.

CSF Update - Journey

- Our journey[3] started last year with an open RFI (request for information), asking if the framework should be updated and, if so, what would they like to see changed?
- We heard from many organizations and affirmed that the framework is still effective today, but it has been five years since a minor update to the framework CSF 1.1, which is the current version.
- There have been a lot of changes in the standards, technology, and risk landscape, such that there are a lot of things that folks wanted to see changed.
- We got more than 100 written comments submissions representing about 4,000 separate comments that we used to make the decision to proceed with the update.
- The RFI was released in February 2022 and followed with four public workshops.
- Our first public workshop attracted more than 4000 attendees from 100 countries around the world. We held several different public comment periods.
- In January 2023, we released an initial draft, called a CSF concept paper[4], that outlined some of the broader themes that we were thinking about changing with the framework. We did an initial discussion draft of the CSF core.
- The current draft CSF Core[5] is the full complete CSF draft that is open for public comment until November 6, 2023. This will be the last draft prior to release of the final CSF in February 2024. We very rarely update the framework, and hence, it is the last opportunity for folks to provide input into the process before the final is released in February.
- Just to provide an update about some of the major changes that we are considering with this new version of the framework.
    - The first one is the scope. Originally, the framework under Executive Order, as well as congressional mandate, is directed at critical infrastructure.
        - Over the past ten years, there has been a much broader use of the framework.
        - Therefore, we decided to remove some of the specific references to critical infrastructure, and just make sure that all organizations, regardless of their sector, and size, are encouraged to use the framework.
        - This means that there are some small changes, e.g., changing the name of the framework. Currently, V1.1[6] is titled "Framework for Improving Critical Infrastructure Cybersecurity, and V2.0[7] will just be the Cybersecurity Framework.
    - There are a couple of other changes throughout the text to reflect its broad use.

Expanding Guidance on Use

[3] Journey to CSF 2.0 https://www.nist.gov/cyberframework/updating-nist-cybersecurity-framework-journey-csf-20
[4] CSF 2.0 Concept Paper https://www.nist.gov/system/files/documents/2023/01/19/CSF_2.0_Concept_Paper_01-18-23.pdf
[5] The NIST Cybersecurity Framework 2.0 https://csrc.nist.gov/pubs/cswp/29/the-nist-cybersecurity-framework-20/ipd
[6] https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf
[7] https://csrc.nist.gov/pubs/cswp/29/the-nist-cybersecurity-framework-20/ipd

- We are also expanding the guidance associated with using framework. This came up quite a bit in the comments.
- Although a lot of organizations are using the framework, they would like to have more guidance on how to use it.
    - We are expanding the guidance around CSF profiles[8], which are tailoring the framework to meet an organization's specific mission and to review their current state versus their target state in terms of improving their cybersecurity posture.
- The biggest change to the framework is the addition of a sixth function: Govern. This was a big deal for us to make this change. The five functions are currently used, and throughout the world as the definition for cybersecurity. They are found in trade policies and standards in legislation, and so it was not a decision that was made lightly.
    - We wanted to make sure that we are reflecting the importance of mission and business considerations within the framework.
    - We have received a lot of praise for the addition of govern. This function highlights the importance of having business considerations in cybersecurity, to have senior leadership set the expectations for cybersecurity within the organization. As illustrated in the CSF wheel with "govern" sitting in the middle, it is those initial governance decisions that will inform the steps an organization will take in response to the other five functions to reduce their cybersecurity risks.
- We also expanded the discussion on third party risks and supply chain cybersecurity.
- We have a whole name category and we ended up including in the new govern function that is focused on supply chain. We also expanded the guidance to make sure that organizations understand how supply chain falls across all the functions of the framework and how to set up their overall supply chain policies.
- Also understand how the framework could potentially be used in discussions with third parties and vendors to oversee their cybersecurity functions as well.

Implementation Examples

- Another big change that we are making with CSF v2.0 is the addition of implementation examples[9].
- We heard there was a gap between the CSF categories and subcategories: cybersecurity outcomes.
- These are higher-level outcomes that we hope an organization will achieve in cybersecurity space versus the informative references, which are the cybersecurity standards and controls.
- Some organizations felt they needed something in between to help understand the intent behind the high-level subcategory. These examples are action-based steps that an organization can take to better understand the CSF subcategories.
- Aimed for use by everybody, but especially smaller businesses will find these useful to better understand how they can achieve some of the subcategories without having to dig up a lot of text and informative references.

NCCoE

[8] NIST Cybersecurity Framework (CSF) 2.0 Reference Tool https://csrc.nist.gov/Projects/cybersecurity-framework/Filters#/csf/filters

[9] The NIST Cybersecurity Framework 2.0 Core with Implementation Examples https://www.nist.gov/system/files/documents/2023/08/07/CSF%202.0%20Core%20with%20Examples%20Discussion%20Draft%5B74%5D.pdf

- We have now run all our workshops, and the last step is the public comment period, which will close on November 6, 2023. Once we receive those comments, adjudicate them and move towards completing the final version as we head into February 2024.
- There is some guidance that is included in the new version of the framework to help with implementation. The guidance has been expanded on how to use the framework, including CSF profiles.
- We have developed and included in the draft a whole new profile template.
- There are other things that NIST and the community can take on to help organizations implement the new version of the framework.
- NCCoE is a collaborative center at NIST, and we work together with industry, academia, and other agencies to solve cybersecurity challenges. We work with various technology vendors to receive hardware and software. We will show how organizations can implement cybersecurity standards and make those standards real using commercially available technologies.
- The other nice thing about NCCoE is we have this really great event space where we host many events. It is about 15 minutes away from the NIST main campus in Rockville, Maryland.

CSF Profiles

- The CSF profiles are a big part of how organizations are going to implement the framework - taking those higher-level CSF categories and subcategories and tailoring them to meet an organization's particular mission.
- Organizations develop profiles for their entire organization and for different business units and different technology lines.
- Community profiles: A trade association or another group will come together and build a kind of sample profile for that community, taking some of the mission specific risks, and the regulations or standards that are subject to that sector and building a profile to do some of the initial heavy lifting for organizations within that sector.
- NCCoE has also published a number of these community profiles. Most of these are done under interagency agreement with other federal agencies, where they are paying us to do this work.
  - We also have executive order mandates and national mandates to develop these kinds of sector specific CSF profiles.
  - We just published our last CSFv.1.1 profile the other day on EV infrastructure.

NIST Publications on CSF Profiles

- There are lots of other nonprofits and trade associations that have developed these community profiles. We encourage organizations to develop profiles for their communities.
- The financial sector profile developed by the Cyber Risk Institute is the most well-known external profile. They are already in the process of updating their profile to follow the new version of CSF.
- Mappings to the framework.
  - The CSF is the central point in which organizations are bringing together multiple standards, regulations, and controls under one framework.
  - Mappings are very important to be able to implement the framework. Additional focus on regulations and mapping regulations and policies, which adds a whole new layer of complexity.
  - At NCCoE, we are also talking about mapping to specific technology capabilities.
  - Spending time on research associated with mappings, trying to define the types of mappings so that we have a common taxonomy to define mappings.

    - Separating out how you can do mappings compared to where to store the mappings in a repository. Will showcase those mappings as part of the NIST cybersecurity program.
  - Also building a new database for the CSF to showcase those mappings.
    - There are different audiences for each of these. There's a broad user set for the CSF, and they are not going to be the ones creating the mappings, but they can use those mappings.

New CSF v2.0 Reference Tool

  - It leverages the Cybersecurity and Privacy Reference Tool (CPRT)
    - Launched this when we launched the new draft and are currently taking feedback on how to improve this tool. This tool is where all the implementation examples are going to live and where all the informative references are going to live.
    - For now, will not keep them in the CSF PDF to allow for us to be able to update them over time.
    - New tool will be helpful for organizations to be able to search by keyword, by different subject terms, and download into different formats to build their own CSF profile.
  - At the NCCoE, we do a lot of mappings,
    - Every one of our 1800 series includes a mapping to CSF and NIST SP 800-53.
    - Those are maintained in PDFs but are reviewing how to convert those mappings to the database. We are working on capturing important information about mappings to specific technology capabilities from these mappings.

Discussions

- **Ms. Flynn Goodwin -** Once you put those mappings in a database, will that be exportable so that agents and organizations will be able to ingest that into their own systems?
  - **Ms. Pascoe -** There are a lot of vendors who have built different tools to be able to ingest this type of information. I can see organizations building on top of this work.
- **Ms. Flynn Goodwin -** This is is a big deal, fantastic.
  - **Ms. Pascoe -** We are excited. It is a big deal because we do not want to break what worked, but at the same time, we really recognized that if we don't do a major update, the CSF would not remain relevant for the future. The CSF has remained relevant for the past decade, and we want the framework to remain relevant for the future decade.
- **Ms. Flynn Goodwin -** Governance right now is interesting because, with the SEC moving forward on publicly traded companies and NIST CSF having been the de facto for so many companies to think about governance for their boards and leadership, making govern more consumable for companies is helpful. Especially since cybersecurity is more of a priority for corporations at that governance tier. I commend you for putting "govern" in and making it more accessible.
  - **Ms. Pascoe -** We received a lot of positive feedback. Initially, we were not quite sure we were going to have consensus on this, but people are already starting to use the draft because they like "govern." We've talked about the importance of business impacts for the framework. It is now a separate function that folks must review, assess, and track. This will make it a bigger priority for organizations going forward.
- **Mr. Scholl –** Still some debate on if it should be inside or outside of the wheel.
  - **Ms. Pascoe -** This might be the only remaining debate. We went through about 20 versions. If somebody else wants to make any changes, they can go through that.

- **Mr. Lipner -** It seems like profiles are very important to organizations that are not big or sophisticated. I heard that other agencies pay you to do them. Are there other small business profiles that are consumable to organizations that don't have the expertise or budget to do a full analysis?
- **Ms. Pascoe -** This new column we are adding on implementation examples in the Small Business CSF profile should be helpful for smaller businesses. There are other organizations that have developed sector specific small business profiles. The rural telecoms and the restaurant industry have developed their own kind of retail profile for those sectors.
  - You do not need a profile to use the CSF, but many organizations find the profile helpful. It is probably more helpful for some of the larger organizations because they are subject to so many different standards and regulations. This is how they are managing everything that they need to do, putting all of it under the framework as their CSF profile.
- **Mr. Gattoni -** How would you communicate to organizations that grabbed onto 800-53 rev4 and rev 5 from a controls perspective, their North Star. Are profiles and mapping the way to facilitate transitions?
  - **Ms. Pascoe –** We hope organizations are not using the framework by itself; that are using it 800-53, ISO 27001, or CIS Top 18 (Critical Security Controls). The framework does a couple of things - it elevates the discussion to a higher level than those standards, making it very helpful to see more strategically where you are today versus where you need to go in the future and to do this gap analysis to improve the five functions, even doing the mapping to the five functions and organizing them. It is very helpful for discussions with senior leadership in understanding what your current cybersecurity posture is. Folks tend to resonate more with the five or six functions than the other ways that these standards are organized.
- **Mr. Lipner –** Is there a mapping to CIS?
  - **Ms. Pascoe -** Yes, we have mappings to CIS Top 18, SP 800-53 R5, and dozens of other standards. Folks can start with the CSF, determine what their priorities are, and then use NIST SP 800-53 and Top 18 to drill down deeper.

Final Board Reviews, Recommendations and Discussions

Steve Lipner, ISPAB Chair

Topics for Future Meetings

- Software liability (ONCD team)
- Data Protection and Vendor Liability
- FISMA Reform (Richard Spires)
- Crypto Agility
- Vulnerability disclosure

Topics brought up on Wednesday by Board Members

- NVD Resources (Ms. Moussouris)
- An issue around the source data and the communication of the known exploitable vulnerabilities. Is it just Feds? Is it at scale? How does it get back to vendors (Ms. Moussouris)
- IoT, IoT data, and IoT privacy (Ms. Moussouris)
- Security in open-source software (Mr. Venables & Ms. Moussouris)
- Barriers to adopting un-biased Machine-Learning Models (Ms. Fanti)

Board Actions

- **Mr. Lipner –** Made a call for actions.

- **Discussion on letter to OMB/DHS/NIST regarding FISMA changes and the CRA**
  - **Ms. Moussouris** – Spoke yesterday about how another letter about CRA would not be helpful. What does the rest of the board think?
    - **Mr. Scholl -** The difficulty is the scope of the board's charter. What would you tell NIST to do about this?
    - **Ms. Moussouris** – I do think that national security concerns are probably the strongest argument there. But I do agree with you.
    - **Ms. Flynn Goodwin -** We have a pivot, which we have not talked about yet. That is the Senate Bill 2251[10], which is the FISMA update and includes provisions around vulnerability disclosure. It's a federal Incident Response System. The bill is buried in the legislation and talks about full disclosure for the federal government. Maybe we want to send some information about that so we can be smarter as a group.
      - The Cybertech Accord group did a paper tying together the FISMA legislation and the CRA; because the principle is essentially the same. Vulnerability disclosure simply for the sake of alerting but not actually doing anything is not helpful.
      - There may be a place for the Board and NIST.
      - NIST was mentioned throughout FISMA reform in the legislation, and so we might want to take up a review of FISMA reform and look at the vulnerability disclosures.
      - This gives us a domestic US and NIST-based connection.
      - Who knows what is going happen when it kicks over to the house. It may go nowhere, and it might die, but that gives us some connection for this issue.
    - **Ms. Miller –** There is a broader effort on the flip to redo FISMA. Richard Spires and his team just delivered something to OMB. So, if you want to look at that, that'll be easier to influence.
    - **Ms. Moussouris –** Is that about writing a letter or getting somebody to talk to us about the FISMA changes at our next meeting?
  - **Decision:** Write the letter and direct it to NIST (OMB and DHS were also proposed. State is out of scope.)

- **Discussion on Crypto Agility**
  - **Mr. Scholl -** I have one more topic. My notes have an idea about improving communication and guidance on crypto agility.
    - **Mr. Lipner –** Volunteered Mr. Venables to do that one. (He was not present at the time.)
    - **Mr. Gantman –** Volunteered to collaborate but not as a primary.
    - **Mr. Lipner -** Usually three or four people are most active in one of these things and then it goes around to the full board.
  - **Decision was made to approach Phil Venables about starting on this, in collaboration with Alex Gantman.**

- **Discussion on Liability**
  - **Mr. Lipner –** Marc Groman had a lot of opinions about liability. I'm not sure that we are far enough along to really say anything about it.

[10] S2251 Federal Information Security Modernization Act of 2023 https://www.congress.gov/bill/118th-congress/senate-bill/2251 introduced 07/26/2023 Senate Committee on Homeland Security and Governmental Affairs. Ordered to be reported with an amendment in the nature of a substitute favorably.

- **Ms. Flynn Goodwin -** There is a provision in the FISMA legislation that would require any entity receiving a federal grant to notify the agency in the event of a breach which would be huge if it does pass. So, if there is a briefing about FISMA reform, that might include some of that, which would go exactly to Marc's point about disclosure of incidents at that academic grantee level.
    - But the FISMA reform is not going anywhere just because of the legislative cycle.
    - If for some reason, there is a convergence, and you start to see the grantee provision then we should talk about it otherwise.
- **Mr. Lipner –** Let's let the vendor liability lie until the next meeting.
    - **Ms. Moussouris –** Agree with waiting but, if we were to sign Marc up for working on it, he was particularly conscious about defining the data protection part of that liability rather than about the security practices.
- **Ms. Flynn Goodwin -** Is there a way to keep in touch with Melanie <Teplinsky>, so that when there is a report out on that issue you can get a copy of it?
    - **Mr. Scholl -** We will follow up after the November 8 workshop.
        - Hopefully the ONCD team that was working on software liability for the National Cyber Implementation plan will have made enough progress to present at the ISPAB March 2024 meeting.
        - The intent was to invite Melanie Teplinsky to set the stage, and then in March when ONCD comes, you will be very informed on what is going on.
    - **Ms. Miller -** That may be a good time to ask Richard to come as well.

- **Decision was made to not write up anything on liability but to invite the ONCD team working on software liability to talk at the March 2024 meeting and Richard Spires regarding FISMA reform.**

- **Discussion on CRA and Conflicts of Laws**
    - **Mr. Gantman –** Regarding the CRA, what if disclosure requires violating laws where you are? For example, if a vendor learned about an exploitation in a classified meeting.
        - **Ms. Flynn Goodwin -** It doesn't matter. Just like it doesn't matter if it's a violation of contract. Conflicts of laws are conflicts of laws, which is why it's such a challenging issue.
            - If you are providing services in Europe, you are required to comply with European law. That is why you saw me pressing her <Christiane Kierketerp de Viron> about national security exemptions.
            - Let's say there was a sensitive issue where you could go to your national competent authority in France, if that were your principal European operations were, and ask ANSSI (https://cyber.gouv.fr/en) to invoke some sort of protection on what is disclosed. There was no such protection.
            - There is this debate, should it be the French CERT or the ENISA team in Athens? The debate is, if it's the ENISA team in Athens, it is a very different prospect and there's no clear indicator that it would go to an intelligence agency, it would go to the national competent authority, which tends to be the CERT.
        - **Mr. Gantman -** Some of the companies that are in the European market have domestic regulations that require them to first disclose to the domestic agency before they disclose to anybody else.
        - **Ms. Flynn Goodwin -** They would have to follow the requirements of the Network Information Systems directive. That would put you with the National competent authority for your principal place of business in Europe, depending upon how it works for you in Europe. It could be Ireland; it could be wherever.
        - **Mr. Gantman –** I didn't mean European manufacturers...

        - **Ms. Flynn Goodwin –** If you are in China, then you will have to disclose to the Chinese government regardless. ‘
        - **Mr. Gantman –** But not in 24 hours.
    - **Mr. Lipner -** It sounds like you just said that there is a conflict between the CRA and the NIS directive.
        - **Ms. Flynn Goodwin -** No, NIS has specific circumstances for reporting an incident, not a vulnerability. But vulnerability reporting is a much bigger deal with financial penalties now for failure to disclose.
        - **Mr. Gattoni -** The CRA reporting requirement is just for exploited vulnerabilities.
- **Mr. Scholl –** I heard a potential communication on vulnerability disclosure concerns and a potential communication on crypto agility.
- **Mr. Lipner -** I will tell Phil what he volunteered for. Happy Holidays!

Next Meeting: The March 20-21, 2024, meeting will be in-person in Washington, DC.

Motion made and seconded to adjourn meeting. The Chair thanked everyone for their participation and adjourned the meeting at 3:30 p.m. ET.

ISPAB – July 12-13, 2023		
Last Name	First Name	Affiliation
Board Members in Attendance		
Lipner	Steve	SAFECode (Chairperson)
Baker	Brett	NARA
Fanti	Giulia	Carnegie Mellon University
Fitzgerald-McKay	Jessica	NSA
Flynn Goodwin	Cristin	Advancing Cyber
Gantman	Alex	Qualcomm
Gattoni	Brian	Federal Reserve Board
Groman	Marc	Privacy Consulting
Miller	Essye	Executive Business Management (EBM), LLC
Moussouris	Katie	Luta Security
Venables	Philip	Google Cloud
Board Members Not in Attendance		
Hallawell	Arabella	WhiteSource
NIST Staff		
Brewer	Jeff	NIST
Scholl	Matt	NIST
Proud-Madruga	Diana	HII/Electrosoft
Elliot	Savannah	HII
Lurie	Kirk	HII
Speakers		
Soutar	Colin	Deloitte & Touche LLP
Megas	Kat	NIST
Teplinsky	Melanie	American University, Washington College of Law
Rajan	Anjana	ONCD/EOP
Djouini	Nasreen	ONCD/EOP
Kierketerp	De Viron	European Commission
Moody	Dustin	NIST
Presman	Dylan	ONCD/EOP
Carpenter	Terry	NSF
Avallone	Linnea	NSF
LaSalle	Connie	NIST

Pascoe	Cheri	NIST
Attendees		
Donelan	Sean	Verizon
Friedman	Sara	IWP News
Ignaszewski	Kate	IBM
Livesay	Jacob	IWP News
Suh	Paul	NIH
Villegas Bravo	Maria	Electronic Information Privacy Center (EPIC)