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 26-27, 2022

Virtual Meeting Platform: BlueJeans

Board Members	**Board Secretariat and NIST Staff**
Steve Lipner, SAFECode, Chair, ISPAB (present)	Matthew Scholl, NIST
Dr. Brett Baker, NARA (present)	Jeff Brewer, NIST
Giulia Fanti, Carnegie Mellon University (present)	Charles H. Romine, NIST
Jessica Fitzgerald-McKay, NSA (absent)	Kevin Stine, NIST
Cristin Flynn Goodwin, Microsoft (remote)	Diana Proud-Madruga, Exeter Government Services LLC
Brian Gattoni, DHS (present)	
Marc Groman, Groman Consulting (present)	
Arabella Hallawell, WhiteSource (absent)	
Douglas Maughan, NSF (absent)	
Essye Miller, Executive Business Management (present)	
Katie Moussouris, Luta Security (present)	
Phil Venables, Google Cloud (present)	

Wednesday, October 26, 2022

Welcome and Opening Remarks

Steve Lipner, Chair, ISPAB, Executive Director, SAFECode

- The Chair opened the meeting at 10:12 a.m. ET and welcomed everyone to the call.
- Introduced new member, Cristin Flynn Goodwin.
- Stated that, while the board provides formal input to the agencies with board letters, the members' comments and questions during the sessions with speakers also add value.
- Reviewed the agenda.

Board Member Introductions and Updates

Mr. Lipner, Chair –

- National Academies of Sciences, Engineering, and Medicine published our report on the future of encryption.
- https://nap.nationalacademies.org/catalog/26168/cryptography-and-the-intelligence-community-the-future-of-encryption

Dr. Baker –

- September 2022: National Archives published white paper on Quantum Information Science and Technology – Implication for Records Management.
- https://www.archives.gov/files/records-mgmt/policy/nara-quantum-information-science-and-technology-whitepaper.pdf
- Looking to strengthen controls and increase assurances in records management.

Ms. Fanti –

- Continuing to study digital currencies and assets as well as synthetic data for data sharing.
- Just finished organizing a conference in Kigali, Rwanda on Digital Public Goods (DPGs), particularly the security, privacy and fairness of DPGs they're implementing.
  - Brought together government officials, NGOs, and academics from across Africa.
  - **https://www.africa.engineering.cmu.edu/research/cylab/summit.html**

Mr. Gattoni –

- It's National Cybersecurity Awareness Month.
  - There is a leadership cadre on the road nationwide to deliver the message of paying attention to cybersecurity.
- DHS/CISA is in partnership with folks at the Israeli National Cybersecurity Directorate.
  - Sponsoring a call for the Israeli-U.S. Binational Industrial Research and Development (BIRD) grant fund.
  - https://www.birdf.com/bird_programs/bird-4/
  - Looking for advancements in AI and machine learning as applied to cybersecurity.
    - Challenging potential providers of services for model portability, model modularity and training datasets solutions.

Mr. Venables –

- Continue to push large-scale pilots of NIST recommended post-quantum cryptography (PQC) algorithms in the Google infrastructure to understand the performance and other characteristics.
- Working on industry-wide product work on software supply chain risk.
  - Providing a leverage point for other organizations to use some of Google's open source and commercial products in that area.
- On the White House President's Council Advisors on Science and Technology (PCAST).
  - Report on Revitalizing the U.S. Semiconductor Ecosystem, September 2022, supported the CHIPS Act.
    - Included recommendations for funding of semiconductor research in supporting crypto and other aspects of security.
  - Will be looking at other activities and potential work on cybersecurity and resilience, probably partnering with NIST.

Ms. Moussouris –

- On three different advisory boards.
- Looking at harmonization of the various new mandatory incident reporting requirements in different sectors.
  - CISA, SEC and FTC reporting, different deadlines, different reporting requirements.
  - Small operators are already stressed in cybersecurity and these differing requirements may not lead to the best cybersecurity outcomes.
  - Maybe invite a speaker to address this in future meetings?

Mr. Groman –

- Teach cybersecurity at Georgetown Law School.
- Look at the diverse range of legal requirement and the inconsistencies between them.
- Focus is on the rapid onslaught of regulatory and legislative proposals around privacy, data protection and data security at the state and federal level across the U.S.

- Trying to help reconcile them and promote consistency.
- Lots of work on privacy risk and harm baked into the NIST privacy framework.

Ms. Flynn Goodwin –

- With Microsoft almost 17 years, entirely in security.
- Also an attorney.
- Focused on advanced attack activity:
  - Nation states, private sector, and offensive actors and ways that we can leverage technology, law, and operational practices to disrupt those types of behaviors from impacting customers and the world.
  - Very busy with the war on Ukraine and watching the evolution of tactics from governments moving away from intelligence collection to destruction and influence operations.
    - Trying to balance technology and what we can do as governments continue to evolve offensive capabilities and Microsoft technology gets stuck in the middle.

Ms. Miller –

- Working with industry partners and agencies across government to address:
  - Continued workforce challenges and how we grow the next generation of cyber talent that fits our needs.
  - DevSecOps and how we can help our industry partners.

Welcome and ITL Update

Charles H. Romine, Director, Information Technology Laboratory (ITL), NIST

Cultivating Trust in IT and Metrology

- **Purpose of ITL is cultivating trust in information technology and metrology**
  - ITL is hosed in the U.S. National Metrology Institute
    - Metrology Institute has a broad set of capabilities:
      - No other Metrology Institute has a counterpart like ITL.

CHIPS Act

- Funding
  - Total funding is approximately $50 billion over 5 years to NIST:
    - Financial Incentives Programs - $39 billion.
      - Intended to catalyze onshoring capabilities for existing chips and forging the next generation of chip manufacturers.
      - Goal to have incentives to bring manufacturing capabilities back onto US shores.
    - Research and Development - $11 billion.
      - National Semiconductor Technology Center (NSTC), Advanced Packaging and Manufacturing Program (APMP), Manufacturing (MFG) USA Institutes, and NIST Metrology research of new capabilities in chip design and manufacturing.
    - Workforce development so we have the workforce necessary to take advantage of these things.
    - Much of the research and development is intended to be executed in grant programs and other kinds of incentives.

CHIPS and Science Act

- CHIPS Act includes appropriations.

- The Science Act is authorized funding and is not yet appropriated.
  - Waiting to see whether appropriations will follow.
- Goals:
  - Supports critical technology research and standards
  - Supports U.S. manufacturing
  - Addresses technology challenges
  - Helps tackle climate change
  - Promotes U.S. competitiveness in international standards
  - Creates a 21st Century NIST
- Intent is to improve the capabilities of NIST over 5 years and includes a potential 40% increase in appropriations for NIST.
- **https://science.house.gov/imo/media/doc/NIST.pdf**

NIST-Wide Critical and Emerging Technologies

- ITL cuts across many different areas of the NIST programmatic priorities
- Will be hearing a lot about these:
  - Artificial Intelligence
  - Biotechnology
  - Cybersecurity and Privacy
  - Advanced Communications
  - Energy Technologies
  - Quantum Information Science
- ITL is the lead laboratory in AI and Cybersecurity and Privacy.
- ITL partners with the Communications Technology Laboratory on advanced communications.
  - Particularly on the security of 5G and beyond.
- Partnered with the Physical Measurement Laboratory on quantum information science.
- Partner with a number of laboratories including the Engineering Laboratory on energy technologies and the Materials Measurement Laboratory on biotechnology.

Zero Trust Architectures (ZTA) – Guidance and Implementation

- Started a ZTA team several years ago at NCCoE.
- Recently published volume B of the preliminary practice for implementing a ZTA.
  - Includes actionable guidance to give people a real understanding of how to implement a ZTA.
- https://www.nccoe.nist.gov/projects/implementing-zero-trust-architecture

Revision to Cybersecurity Framework (CSF)

- Updating CSF to version 2.0.
- Seeking input from the private sector community.

DevSecOps Workshop

- This is important to ITL.
- Looking at better integration of software development and security.
- Making development of software more agile and shortening the life cycles.
- Incorporating newer technologies rather than ignoring issues of things like cloud computing and taking advantage of the newer technology practices.

Artificial intelligence update

- ITL magnifies its impact through collaborations with the private sector.
- Developed a risk management framework for AI.
  - Held workshops – most recent was on October 18-19.
  - Used our experience in the identification of AI bias and mitigating that bias.
  - Goal is to accompany the risk management framework in the context of a risk management framework playbook.
    - Playbook will give additional guidance on actionability. We'll make it a more dynamic environment so we can continue to progress AI risk management.

National Initiative on Cybersecurity Education (NICE)

- Working on workforce development.
- Engaged with the private sector to identify where the needs are.
- Entered into a 5-year cooperative agreement with Katzcy on cybersecurity games.
  - Designed to develop needed skills.

Congressional Hearings or Briefings

- Three opportunities to acquaint elected officials in the House Committee on Science, Space and Technology, Subcommittee on Research & Technology with the work ITL is doing:
  - September 29, 2022: Trustworthy AI: Managing the Risks of AI
  - July 28, 2022: Exploring Cyber Space: Cybersecurity Issues for Civil and Commercial Space Systems
  - June 29, 2022: Privacy in the Age of Biometrics

Celebrations, Recognitions, and Other Miscellaneous Information

- NIST is working on balancing the return to on-site work with continued remote work.
- Celebrating 75 years of applied math and statistics and 50 years of cybersecurity research at NIST.
- In 2023, NIST will celebrate 60 years of Biometrics research.
- September 29th, ITL launched the NIST Cybersecurity Program History and Timeline (https://csrc.nist.gov/nist-cyber-history)
- Recognitions:
  - Ram Sriram: achieved the status of IEEE Life Fellow Status in recognition of many years of loyal membership and support of the activities.
  - Raghu Kacker: recently named Fellow of the Washington Academy of Science for outstanding contributions in enabling the field of combinatorial testing in measurement science and software engineering.
  - Kamran Sayrafian: recently named Fellow of the Washington Academy of Science for recognition of outstanding contributions to mathematical and computational modeling of Body Area Networks.
  - Paul Patrone: received the Excellence in Research in Applied Mathematics Award from the Washington Academy of Science.
  - Barbara Guttman (SSD): received the Leadership in Forensic Science Award from the Washington Academy of Science.
  - Members of the Image Group in IAD were recognized with the Service and Leadership Award at the 2022 Federal Identity Forum.
  - The Image Group's contactless fingerprint research team was also recognized as a finalist for the 2022 Federal Identity Forum Best Educational Effort Award

Discussion

Question on the adjacency of ITL and metrology

- Mr. Venables: There is an absence of great work on measurement in security and privacy. There is an industry challenge to define metrics and measurements for the right outcome.
- Mr. Romine:
  - Vint Cerf, the father of the Internet, is an advisor to ITL. He has been pushing for us to:
    - take some of the metrology principles that NIST is known for and apply them to cybersecurity and privacy.
    - Determine which metrics would be useful
  - Implications are that behavior tends to follow what you measure. Need to be careful in the choice of metrology.
  - Currently conducting research in this area.

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

AI Risk Management Framework

Elham Tabassi, NIST

NIST AI Program

- AI is transforming the world.
- NIST is looking at cultivating trust in AI by:
  - Conducting foundational research to advance trustworthy AI technologies,
  - Advancing AI research and innovation across NIST's laboratory programs,
  - Establishing benchmarks and developing metrics to evaluate AI technologies,
  - Participating and leading in developing standards to advance AI innovation,
  - Contributing NIST's technical expertise to discussions and development of policies,
  - Ensuring NIST has resources and expertise to carry out its AI programs

Key NIST Roles for the Federal Government

- NIST AI RMF is a congressional mandate as part of the National AI Initiative Act of 2020
  - See Division E of https://www.congress.gov/116/bills/hr6395/BILLS-116hr6395enr.pdf
- National AI Advisory Committee (NAIAC)
  - https://www.ai.gov/naiac/
  - Established an office at the Office of Science and Technology (OSTP)
  - NAIAC established April 2022.
  - Held their first meetings May 4 and October 12-13, 2022.
  - Tasked with advising the President and the National AI Initiative Office on topics related to the National AI Initiative (https://www.ai.gov/).
- AI Research Resource Task Force
  - https://www.nsf.gov/cise/national-ai.jsp
  - Also established by the National AI Initiative Act
  - Task force:
    - Consults with experts and stakeholders from government agencies, private industry, academia, and civil and disabilities rights organizations
    - Informed by ongoing interagency efforts

    - Leveraging cloud computing resources in support of federally funded AI research and development
- Federal AI Standards Coordinator
  - https://www.nist.gov/standardsgov/icsp-ai-standards-coordination-working-group-aiscwg-charter
  - Purpose is to facilitate the coordination of federal government agency activities related to the development and use of AI standards and develop recommendations relating to AI standards.
- Interagency Coordination:
  - White House OSTP & National Science and Technology Council (NSTC),
- Stakeholder outreach is the basis of everything we do.

Trustworthy and Responsible AI @ NIST

- Focus is on a trustworthy, responsible AI program that has the purpose and goal of cultivating trust in the design, development, and use of AI systems.
- We need to define what we mean by "trustworthy".

AI RMF: Management Risk Through Trustworthy and Responsible AI

- Trying to address risks to individuals, communities, organizations, and society.
  - Looking at maximizing positive and minimizing negative impacts
- This is a congressionally mandated, living document for voluntary use
- Aimed at preserving rights and operationalizing values
  - Non-discriminatory and understandable
  - See Executive Order 13960 https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government and Organization for Economic Co-operation and Development (OECD) AI recommendations https://oecd.ai/en/ai-principles
- Is law and regulation agnostic
- This will be achieved through an open, transparent, collaborative development process that is outcome based, secure, private, and enforceable by design.
- Need to think about AI system profiles in context
- Have an interoperable lexicon to communicate about risk with a foundational setting for assessing trustworthiness

AI RMF Timeline and Engagements

- Started with an RFI in July 2021 and workshops
  - Results are posted on the AI RMF website: https://www.nist.gov/system/files/documents/2021/10/15/AI%20RMF_RFI%20Summary%20Report.pdf
  - Recordings of workshops: https://www.nist.gov/itl/ai-risk-management-framework/ai-risk-management-framework-workshops-events
- Used input from RFI and workshops to create an AI RMF concept paper for public comment
  - Published December 13, 2021
- Will repeat the process of draft document, public comments, listening sessions, and workshops 3 times with the third workshop on October 18, 2022.
- Intend to publish version 1.0 in January 2023.
- Also published research papers in support of different concepts
  - Explainable AI paper released September 29, 2021

  - Bias in AI released March 14, 2022
- Have reached out to industry NGOs, academia, civil society, and standards organizations for the listening sessions

Transforming Culture – Socio-technical Systems Approach

- Takes into consideration the larger social context in which AI operates, its purpose and potential impacts
- Trying to transform the culture and establish a socio-technical approach to understanding AI risks
- AI systems have complex interactions between the system's environment and humans.
  - All these aspects need to be understood
  - AI systems are all about the context of use.
    - Using risk management to understand risk of things like face recognition
- Should take a human-centered design approach
- Need to apply the scientific method to AI systems
  - Ensuring proper validation and verification prior to releasing code
- Set up governance structures for the people who build and maintain AI systems
  - Management is a shared responsibility across everybody involved at every stage of the AI lifecycle.
  - Having the policy and procedures in place is important
  - The role of the senior leaders and senior management is important to setting the culture of understanding risk and risk management and ensuring organizational buy-in.
  - Don't want to rely on a checklist for compliance
- Include consideration of limitations from an impact and values-based perspective.

Trustworthy AI Characteristics

- Goal is that trustworthy AI systems achieve a high degree of control over risk while retaining a high level of performance quality.
  - Requires a comprehensive approach to risk management, with tradeoffs among the trustworthiness characteristics
- All characteristics fall under an accountable and transparent umbrella
- Characteristics include:
  - Valid and Reliable
  - Safe
  - Fair and Bias is managed
  - Secure and resilient to different vulnerabilities
  - Explainable & interpretable
  - Privacy-enhanced
- Many hours have been spent discussing each of these characteristics with the community and amongst ourselves.
- We are building research programs for each characteristic.
- The community has also built tools for measuring each characteristic, evaluating platforms, and benchmarking systems.
- These characteristics are interdependent
  - It's impossible to maximize all. There will be trade-offs.
  - In different contexts some characteristics will be more important than others.

Question from Ms. Flynn Goodwin:

- Those who are at the cutting edge of developing and releasing AI projects may not be consuming the nuances of risk management. Just as there is a minimally viable product, is there a conversation around a minimally acceptable risk?
- Ms. Tabassi reiterated that the AI RMF is intended to be context agnostic. The minimum amount of safety should be part of the discussions of the AI profile for that vertical. It's the job of the community that's going to use it to define these things. Our next step is to get the community engaged.
- Ms. Moussouris asked about the risk of amplifying the bias that exists in systems and the measures being taken to measure the disparate impact on certain groups.
- Ms. Tabassi replied that this is a difficult thing to do. That's why we need to take a socio-economic technical approach. Most of the talks of bias refer to statistical and computational bias but human and systemic biases also need to be understood. We collect data from the environment and society to build algorithms. The bias is in our society and that gets into the training sets that are going to be part of the system. The Bias in AI paper addresses this (https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf). We are working on understanding how to measure bias and each of the trustworthy characteristics and having a terminology taxonomy. We are getting the community together to build evaluations around that. The NCCoE has a project looking at bias in the context of credit underwriting. We haven't solved bias, but we have a way forward and we have community engagement to get there.

AI Risk Management Framework Core

- Follows the structure of the Cybersecurity Framework and the Privacy Framework.
- Provides guidance on how to map, measure, manage, and govern AI systems. These functions aren't performed sequentially. There can be a lot of back and forth.
  - The Map function is about understanding the context and the trade-offs as discussed previously
  - The Measure function is about measuring the identified risks
  - The Manage function covers guidance on how to prioritize the risks and the right level of resources based on the magnitude of risk and the impact.
  - The Govern function deals with how the risk management culture is cultivated and present. This includes the rules, responsibilities, culture, and policy and procedures

NIST AI Risk Management Framework Playbook

- This is an online companion resource for the AI RMF.
- Includes suggested actions, references, and documentation guidance for stakeholders to achieve the outcomes for the map and govern functions.
- The other functionalities will be added in the future.

Question from Mr. Groman:

- In the current documents, is there an inventory, catalog, or library of potential harms to be considered so that the risk framework references potential harms to individuals, the enterprise, and society? The documents refer broadly to civil liberties considerations, which may not mean a lot to the average engineer or designer, but in some cases the negative impact of bias could be serious. For example, in the police context, an adverse or negative outcome from a biased AI algorithm could be incarceration while in another context it could be much more mundane.
- Ms. Tabassi replied that OECD works closely with national and international partners on policy and building quality assurance. There are also documents coming out of CISA. When working on the AI RMF profiles and evaluations, understanding a risk catalog will be part of the discussions. The Bias

in AI document does look at some of the negative impacts of bias in specific use cases. We will consider adding more to the playbook.

AI RMF Profiles

- Use-case profiles: Instantiations of the AI RMF functions, categories, and subcategories for a certain application or use case based on the requirements, risk tolerance, and resources of the Framework user.
- Temporal profiles: descriptions of either the current state or the desired, target state of specific AI risk management activities within a given sector, industry, organization, or application context.
- NIST welcomes contributions towards the development of AI RMF use case profiles as well as current and target profiles.

Crosswalks

- There is no shortage of incredible documents, and we started doing a crosswalk between the AI RMF, OECD AI Recommendations, the proposed EU AI Act, and EO 13968.
- In discussions on development of the crosswalk with a Singapore AI governance model and ISO documents.

NIST AI RMF Related Resources

- AI RMF Profiles
- AI RMF Playbook
- AI RMF Glossary
  - Bring all the different terms together
- AI Standards Hub
  - UK just released AI standards
  - This hub will track standards and training materials and provide a place for engagement.
- AI Metrics Hub

Discussion

Guidance on potential negative impacts

- The Chair reiterated the importance of providing developers with a resource that identifies a set of bad things that an AI application really ought not do. He indicated that he sees the RMF as the "keystone" product out of NIST in AI privacy and security and it would be a missed opportunity if it doesn't provide clear guidance.

Measuring Fairness

- Mr. Scholl read a question from a board member (didn't hear name). Does the framework help organizations evaluate how they should measure fairness?
- Ms. Tabassi replied, the reason they're looking at Resource Center metrics is to put the multiple dimensions for fairness and other characteristics together, get the community to look at that, put together practices in getting people to use the different metrics and report of the use of the different metrics; what works and what doesn't work for fairness. That's also the case for the other metrics as well as terminology. Evaluations of AI are important, looking at what is monitored, and getting input on if it worked or not, where it could be improved and how. Is it meeting stated objectives and what are the outcomes as well as being able to measure progress toward those outcomes and engaging with the community on how to measure effectiveness that's important?

The Chair recessed the meeting for a 45-minute lunch break.

NIST Frameworks: Unpacking and Settling with 3 NIST Frameworks

Dave Weitzel, MITRE
Julie Snyder, MITRE
Christina Sames, MITRE

Introductions

- Speakers introduced themselves and gave some background on MITRE.
- Working on clarifying how the NIST Cybersecurity, Privacy, and Risk Management Frameworks should be used together so they:
  - Complement each other,
  - Provide distinct benefits, and
  - Result in collective power.

Overview of the Three NIST Frameworks (covered very briefly)

- Typically start sessions with an overview of the frameworks.
- There are always folks new to the process.
- Talk about origins, the timing, where they are voluntary and where they aren't, what the purposes are, how they were developed, etc.
- The three frameworks:
  - Cybersecurity Framework (CSF)
  - Privacy Framework (PF)
  - Risk Management Framework (RMF)

Cybersecurity and Privacy Framework Components (covered very briefly)

- They then talk about the CSF and PF together since they are similarly structured.
- Core
  - Increasingly granular set of activities and outcomes that enable organizational dialogue about managing risk
- Profiles
  - Prioritized subset of the Core that addresses risk in alignment with organizational objectives
- Implementation Tiers
  - Helps an organization determine if it has sufficient risk management practices and resources in place to achieve its Target Profiles.
- Focus mainly on profiles and core in their presentation.
- They talk about the privacy framework and establishing or improving a privacy program.
- Go into more of the technical details of products, systems, and services for privacy risk management.
- Emphasize that you don't have to be an expert to understand the functions in each of the frameworks.

Cybersecurity and Privacy Framework Cores

- Introduce the cores from both frameworks
- Each framework has five functions.
- There is general alignment at the top level between the functions
- The PF points to the CSF for detect, respond, and recover

- The common terminology of these frameworks aids in communicating risk across an organization, with management, and when talking to the board of directors.
- Talk about the value of having the frameworks as a starting point for organizations that are new or might not yet be achieving an activity or outcome.
- Frameworks point out activities organizations can use to map their own policies and standards

Risk Management Framework (RMF)

- Has gone through a recent update
- Identify where it can complement the CSF.
- Highlight activities that organizations can do to help them be ready to implement the RMF.
- Creating a connection with the CSF for a strategic type of framework that isn't system specific, or system focused.

RMF Steps

- Mainly focusing on what the Prepare step is intended to do.
- These steps are not sequential

Framework Profiles

- Profiles mark the way that you tune the CSF and the PF for use by an agency, enterprise or organization.
- You look at the organization, what it's trying to do, what legal framework it's in, its requirements, and how it structures itself.
- To do that, you take the core and adapt it
    - Instead of having all the components of the CSF, you tune it to your enterprise and focus on those functions, categories, and subcategories early on to help address risk gaps in the enterprise.
        - Now map to your current state. This gives you your current profile and target profile
- For many industries, you can start with a notional target profile and adapt it to your own enterprise needs.

Understanding How to Draft Mission Objectives: Terminology

- The mission is a CEO's or agency head's elevator speech
- The top of a profile is a mission objective.
- Business/Mission Objectives: The fundamental purposes and operations of an industry/subsector or organization that the processes and systems support.
- Mission Priority: The relative importance of one item versus another. Which objectives have highest priorities for keeping the organization running and minimizing the risk.
- Mission Dependency: A requirement to fulfill Mission or a Mission Objective that lives outside of the subsector.

Example: Maritime Bulk Liquids Transfer (MBLT) Profile

- Some mission objectives:
    - pass required audits and inspection,
    - maintain quality of product
    - obtain timely vessel clearance
    - maintain personnel safety
- Often, two of the top mission objectives are keep people safe and don't damage the environment

- Once you have all the mission objectives,
  - They need to be prioritized into a list and a process
  - They then get mapped against the five functions:
    - Identify, protect, detect, respond, and recover
  - For this CS profile, there are 23 categories.
    - Each category breaks out into subcategories.
  - For each category and subcategory, the objectives are prioritized using a dot chart (the numbered columns refer to the 8 identified mission objectives):

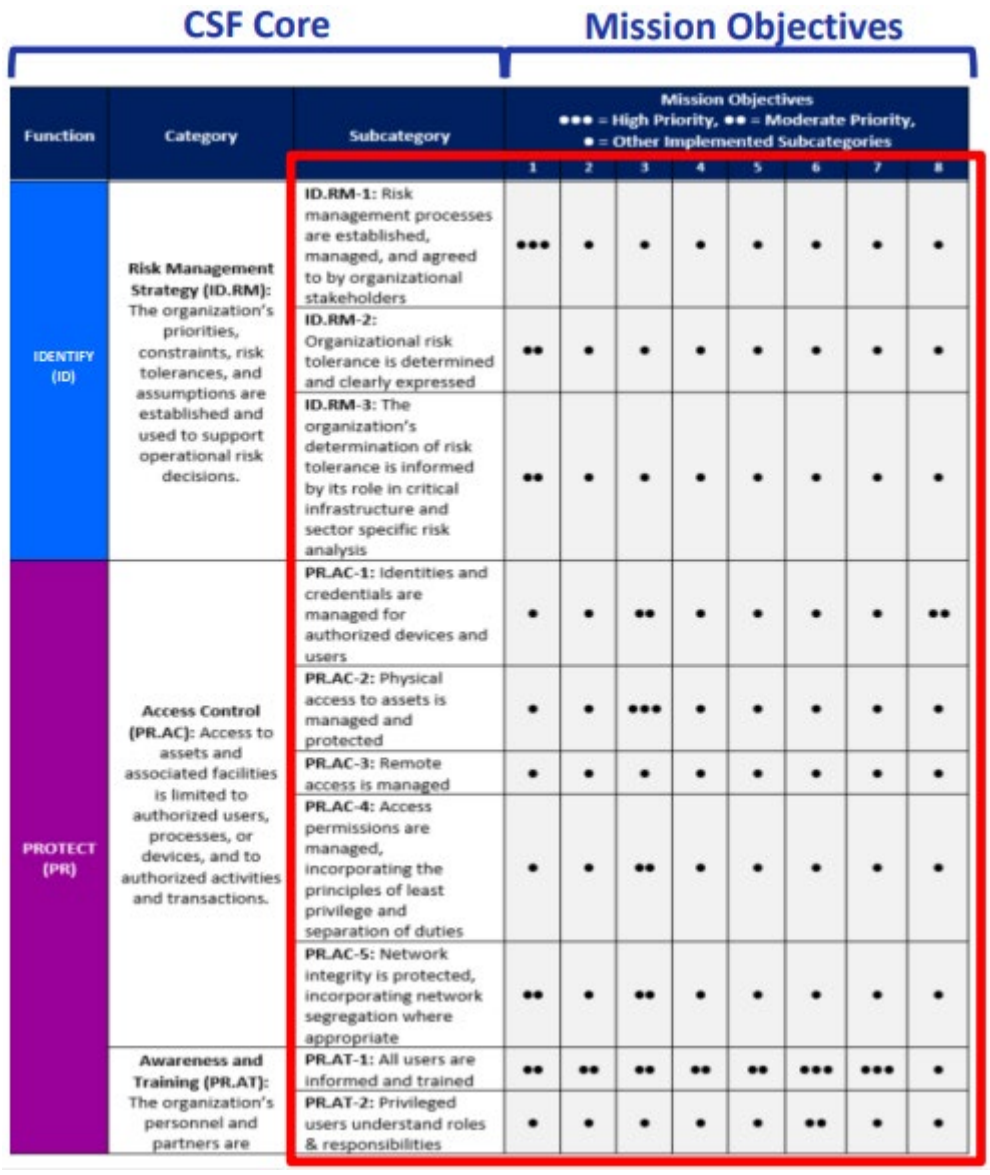

CSF Core | Mission Objectives

Function	Category	Subcategory	Mission Objectives ●●● = High Priority, ●● = Moderate Priority, ● = Other Implemented Subcategories							
			1	2	3	4	5	6	7	8
IDENTIFY (ID)	**Risk Management Strategy (ID.RM):** The organization's priorities, constraints, risk tolerances, and assumptions are established and used to support operational risk decisions.	**ID.RM-1:** Risk management processes are established, managed, and agreed to by organizational stakeholders	●●●	●	●	●	●	●	●	●
		ID.RM-2: Organizational risk tolerance is determined and clearly expressed	●●	●	●	●	●	●	●	●
		ID.RM-3: The organization's determination of risk tolerance is informed by its role in critical infrastructure and sector specific risk analysis	●●	●	●	●	●	●	●	●
PROTECT (PR)	**Access Control (PR.AC):** Access to assets and associated facilities is limited to authorized users, processes, or devices, and to authorized activities and transactions.	**PR.AC-1:** Identities and credentials are managed for authorized devices and users	●	●	●●	●	●	●	●	●●
		PR.AC-2: Physical access to assets is managed and protected	●	●	●●●	●	●	●	●	●
		PR.AC-3: Remote access is managed	●	●	●	●	●	●	●	●
		PR.AC-4: Access permissions are managed, incorporating the principles of least privilege and separation of duties	●	●	●●	●	●	●	●	●
		PR.AC-5: Network integrity is protected, incorporating network segregation where appropriate	●●	●	●●	●	●	●	●	●
	Awareness and Training (PR.AT): The organization's personnel and partners are	**PR.AT-1:** All users are informed and trained	●●	●●	●●	●●	●●	●●●	●●●	●
		PR.AT-2: Privileged users understand roles & responsibilities	●	●	●	●	●	●●	●	●

- This is useable by executives to talk about prioritizations and ensure funding is allocated accordingly.

Observation from Mr. Groman on "Privacy" definition

- It is likely that even the experts in this room do not know how NIST has defined privacy in the privacy framework and what it means to assess privacy risks and privacy harms or problematic data actions. "Privacy" isn't intuitive, implementable, or actionable. It's problematic.

- Explaining what the problematic data is, the action is, and the associated risks would help distinguish between the actual privacy risks around bad things happening from the authorized use of data versus bad things happening from the unauthorized use of the data.
- Ms. Snyder commented that most of the conversations on privacy risk management engineering are in NISTIR 8062. The term "privacy" is open ended for people to apply it in a way that suits their view of the universe when it comes to privacy risks. NIST has a Venn diagram that shows the relationship between cybersecurity risk and privacy risks. In the overlap of the diagram is protecting personal information (PII). Part of that notion is, from a security perspective and protective privacy perspective is about the unauthorized use and access of data. The rest of the privacy circle in the diagram considers what else might go wrong for privacy, even when you're doing something that's a valid business or mission purpose.
  When NIST was researching terminology, many of the terms used for privacy got translated in people's heads into security functions, missing the privacy nuance that gets to those things that go wrong during authorized processing.
  The term "problematic data actions" is data flows that might introduce privacy risk; looking for data flows where there might be some privacy risks.
- Mr. Groman clarified that the word "privacy risk" is meaningless to someone building a system. They're all bringing their own definition of privacy.
- Mr. Venables added that in the privacy framework there are different roles in an organization that must pay attention to different things. The engineers don't really think about privacy. They think about controls they have to implement. Some privacy properties may be specified by a product manager and the product manager has to think about those authorized uses that would cause a problem. Does the framework differentiate which organizational roles have to think about which aspect?
- Ms. Snyder replied that it acknowledges that you have different roles, and you need to understand them, but it doesn't get to that level of detail. Many organizational privacy programs can sit in a variety of different places such as: legal, under the CIO, or as their own program. That makes it hard to establish rules that work all the time, but it does give you a construct to start thinking about those things.
- Mr. Venables added that this might be something useful to think about. Some of the biggest privacy issues in industry, where there were gaps, are where the privacy lawyers thought the engineers were implementing certain controls or the engineers built a perfectly specified set of controls, but still led to privacy issues because the product itself was not sufficiently thought through or public expectations changed, thus creating a privacy issue.
- Mr. Wetzel mentioned a conversation he had with the Dutch data protection minister. The Dutch minister said, "You Americans are so funny. You talk about privacy being a constitutional right, but you never do anything about it. In Europe, we declare it to be a right and we get on about it, but we call it data protection. Because of that, it moves right into the protection discussion."
  When Mr. Wetzel talks with cyber engineers who don't quite understand privacy, he brings up data protection and they understand much better.

Ways to Use a Profile

- Customization of the Core for a given sector, subsector, or organization
- Fusion of business/mission logic and cybersecurity outcomes
- Alignment of cybersecurity requirements with operational methodologies
- Basis for assessment and expressing target state (where you want to go)

- Decision support tool for cybersecurity risk management
  - Can be communicated up through management into C suites.

Framework Profile Development Process

- They mentioned that there is a process that is very thorough but did not go into detail in this presentation.
- Pointed out that the third step from the bottom, Create Profile Document, is often seen as the end of the process but there is still governance and allocating resources and budgeting.
  - This speaks to the idea that if you have extra money, this provides guidance on where to allocate that money.

Ms. Moussouris asked about vulnerability disclosure programs

- She asked how they are treating preparation, and the allocation of resources, for ongoing support of vulnerability disclosure programs as described in the NIST framework? This is about responding to vulnerability reports from researchers.
- Mr. Wetzel declined to answer because he didn't have any vulnerability disclosure people with him. He did indicate that he would talk with them and get back to Ms. Moussouris.

Risk Management Framework (RMF)

- In their workshops, they present the basics of the RMF to ensure the audience has a foundational understanding of what goes into the major components of the RMF including the major documents that support implementation of the different steps.
- They then do a walk-through of what a NIST SP 800-53 control looks like.
  - Highlight the different components and parameters that might affect the implementation based on organizational needs.
- They also present a NIST baseline example. In this case they presented the baseline table for the Awareness and Training (AT) family of controls
  - The example highlighted that there's the new privacy control baseline
    - This is to address some potential communication gaps and risk management activities that could have impacts from either a security or a privacy perspective.
- They then cover what an overlay is. Included for each control in the overlay is:
  - Justification for inclusion
    - May be legal or regulatory reasons
  - Control specifications (are there protections that need to be in place):
    - A selection indicator (+, -, or blank)
    - Guidance
    - Parameter Value
    - Control Extensions
    - Reference to the applicable requirements

How Framework Profiles May Influence the RMF Steps

- Profiles can help an organization communicate what needs to be done when implementing the RMF.
  - Identifies:
    - priorities for cybersecurity and privacy activities
    - desired outcomes
  - This helps identify necessary controls and what kinds of monitoring might be required
- Profiles can also help identify what resources are needed and where

How the CSF and PF can help RMF activities

- Helps with the communication
  - Not everyone understands controls and the technical terms used in the RMF
  - Using the CSF categories, sub-categories, functions and other foundational CSF and PF concepts allows technical experts to have conversations with leadership
    - Helps create the connection between what's being done at a system level and what strategic priorities leadership is seeing to support their mission.

Bringing the Frameworks Together

- They did a presentation for RSA.
  - 45-minute presentation followed by a 3-hour hands-on interactive lab where they walked through these processes using a hypothetical system.
  - The hypothetical system RMF baseline that was developed had 450 controls that needed to be implemented. This was overwhelming.
  - They then showed how, if they have a target profile from the CSF and/or PF, that can be used to help prioritize and otherwise make sense of the RMF baseline.
    - Mapped the mission objective priorities from the profile to the RMF baseline controls.
    - This also helped identify gaps in the baseline controls, identifying controls that could be added during tailoring.
    - Of the 450 controls, 12 controls were identified as having a high priority.
    - Other resources were also identified that could support implementation.

Communications to Manage Risk Effectively

- These frameworks can work together to help promote communications about risk and risk management throughout an organization, including with entities not directly involved with security or privacy.
- Risk management is not a one-size fits all approach. There is flexibility within all three frameworks to work together and to manage risk.
- Using the words within a CSF/PF to:
  - understand what's being implemented at the system level,
  - Make implementation decisions about what's being done on a system,
  - Have that information go back up to leadership to understand the threats and vulnerabilities to systems enabling them to make strategic decisions on possible changes to a risk posture or resource allocations.
- **Discussion**

Question from Mr. Venables about automation

- He asked what they think of emerging standards like OSCAL to turn this into something machine readable and create a machine-driven systematic process, especially once it has been implemented and they want to sustain it. They want to monitor if the controls are effective.
- Mr. Wetzel reiterated that Mr. Venable is asking if this can be automated once it's in a steady state. He mentioned that NCCoE has had a demonstration lab with DHS for a long time and he recommended contacting and talking with them because it's critical to be able to automate portions of this. He doesn't think it will be possible to automate everything. Sometimes there are going to be conflicts that have to be balanced and automation would not work.

Ms. Miller asked about where these discussions are taking place and how do we expand them?

- How do we get this conversation to the CIO Council, CISA Council, Privacy council, etc.…?
- Mr. Wetzel replied that they did a proto profile for a cabinet level agency, and it was hard to get above the CISO level due to lack of a governance structure. There are three aspects, your technical component, what are you doing with budget and resource allocation, and how do you govern it.
- Ms. Snyder added that when they were in the thick of working with the oil and gas industry using the Coast Guard profiles, one of the major players had a metal plaque with the five functions of the CSF.
- Mr. Wetzel added that one of the chief cyber architects of a major oil company indicated that leadership was coming to him after a major event asking him to come to their meeting and explain the risk governance process in terms they could understand. He was able to have those discussions. He confirmed, however, that there is a gap. Many enterprises haven't crossed the divide.
- Ms. Snyder also added that NIST is working on making sure that the terms and processes from different publications work together. Having conversations looking at the same thing from different perspectives.
- Mr. Groman said that NIST is the resource, but NIST is not a regulatory authority and has no authority to make this binding. That falls on OMB. They oversee those councils. The OMB Circular A-130 and others do make it mandatory for federal agencies to have to comply with 800-53 and the other documents. But there are so many competing requirements for agencies and IT staff. He believes there has to be enhanced communication.
  He mentioned that there has been a mass departure of privacy professionals over the last 10 years that is slowly being filled. He suggested that, instead of talking about "privacy", let's talk about bad stuff that can happen when you process data. He recommends moving away from ambiguous terms and talking about what happens to facilitate these conversations between various people like CTOs and CISOs.
  The processes in these documents are not self-executing. Context is important. The same collection of data by one agency does not present the same risk when done by another agency.

The Chair asked about the low, moderate, high baselines in 800-53.

- He mentioned that it seems that the controls in 800-53 are prescribed, based on the system sensitivity determination. He was on a committee that met with the IT staff of an agency at the operations level. These people were starting implementation of RMF and trying to meet the requirements. Their understanding was that they couldn't manage their systems securely because their priority was on the RMF requirements.
- Ms. Snyder agreed that is how some use 800-53 baselines but that doesn't take into consideration the uniqueness of that system; what it's inheriting from other systems, the operating environment, or the users that have access to it. The baseline is a starting point and then you're supposed to tailor and bring in controls where you need to and eliminate those you don't need. That's where overlays can help. There's a privacy overlay available, there's a classified systems overlay, etc. These can help with tailoring.
  This is also where communication is important. Understanding what you're doing with the system and what purpose it's designed to provide. That was also why we created the working example with 450 controls; to show the need for prioritization and providing guidance for implementing controls as opposed to just having a checklist of controls that you work through. Creating a checklist is easier at the outset but the example shows how a CSF or PF profile can help prioritize which controls are most important.

Mr. Groman talked about residual risk.

- He mentioned that when he was a senior advisor for the Privacy Council, he had to try and translate the idea of owning their residual risk. He's been in a position where a cabinet secretary was insisting that he sign-off on something and he had to point out that it's not his place to do that. His job was to look at and assess the system and identify privacy risks or risks to individuals from the system and then to assess the severity of it and provide you with ways to mitigate that risk. If even after we mitigate that risk there's still going to be residual risk, you have to decide if you're willing to own it. There are three options: the system owner could decide the mission is so important the risk doesn't matter, and they will go forward, the risks identified are so bad the system owner kills the program, or the risks are able to be mitigated sufficiently and the system owner is willing to own any residual risk. We don't want to get to zero risk. That's called no data and not working. In privacy and security there is always risk. He wants people to work toward the idea of owning the risk.
- Ms. Snyder added that this is where profiles can be a helpful tool. Profiles connect your enterprise view of mission and business objectives with cybersecurity and privacy activities. Instead of doing cybersecurity or privacy just to meet the regulation and be compliant, you know what the activities are supporting.
- The Chair commented that it's important to get beyond the abstract control catalog model to actionable packages of controls that enable agencies to operate securely at low cost. He thought he saw something recently from CISA about prepackaged control checklists that would help agencies avoid having to go into every setting.
- Mr. Wetzel commented that for some entities, checklists would be an improvement over the use of point solutions where they are buying and implementing tools on a crisis-by-crisis basis. As a result, they have implemented the same control in multiple ways, many of which are paid for but not being used. If they don't have integration between, at the very least a checklist, governance, and budget, then there's a lot of room for improvement.

The Chair recessed the meeting for a 5-minute break.

Entropy Sources: Importance and Testing

Meltem Sonmez Turan, NIST
John Kelsey, NIST
Tim Hall, NIST

Part 1: Overview of NIST Standards on Random Bit Generation

- Cryptographic Random Number Generation
  - The security of cryptographic primitives relies on the assumption that bits are generated uniformly at random and are unpredictable.
    - Typically used to generate cryptographic keys and nonces.
  - Designing random bit generators (RBGs) is challenging
    - First challenge is finding robust random sources and correctly extracting randomness
    - Also, difficult to understand how unpredictable these outputs are.
      - This corresponds to the problem of estimating entropy
    - Difficult to statistically model the process
    - Difficult to understand the effects of outside parameters and environmental conditions on the source.
      - We're also in an adversarial setting where people might try to guess these numbers which adds to the difficulty in understanding these effects.
  - Validating RBGs is also challenging

        - Need some expert knowledge on the random sources themselves
        - It's difficult to verify and understand some of the claims
        - There are also practical constraints (e.g., time)
            - We want this process to be efficient
- NIST SP 800-90 Series
    - Provides guidelines on how to construct RBGs that are validated through FIPS 140.
    - Aims to improve the quality of RBGs by specifying design principles and requirements.
    - Has three parts:
        - 800-90A: has recommendations for random number generation using deterministic RBGs (DRBGs)
        - 800-90B: has recommendations for the entropy sources used for random bit generation
        - 800-90C: has recommendations for RBG constructions
            - 90C is about combining 90A and 90B
- NIST SP 800-90A
    - Specifies mechanism for the generation of random bits using based on hash functions and block cyphers
    - The original version came out in June 2006.
    - It was updated in January 2012 and NIST is working on a new revision to align the requirements, description, and conditions with the new 800-90C
- NIST SP 800-90B
    - Provides an entropy source definition and model
    - Specifies design principles and requirements for entropy source components
    - Includes entropy estimation techniques
    - The original version came out in August 2012
    - It was updated January 2016 and NIST is planning to revise it based on lessons learned during validation testing
- NIST SP 800-90C
    - Describes the three RBG constructions:
        - The first provides random bits from a device that is initialized from an external RBG
        - The second includes an entropy source that is available on demand
        - The third includes an entropy source that is continuously accessed to provide output with full entropy
    - Initially published in August 2012
    - Updated April 2016
    - A third public draft was published September 2022 with public comments due December 7, 2022
- NIST SP 800-22: Statistical test suite for random and pseudorandom number generators for cryptographic applications
    - This specifies 15 statistical randomness tests, including a software tool
        - Can use these tests to analyze your random number generator
    - Originally published in October 2000
    - Updated August 2008
    - The crypto publication review board recently reviewed it and proposed revisions to the standard to align with SP 800-90 series and to make technical improvements
    - NIST is working on revision 2
- Aligning NIST and BSI standards
    - BSI in Germany is similar to NIST, and they also have standards on random number generation

        - AIS 20: is for functionality classes and evaluation methodology for deterministic random number generators
        - AIS 31: is for functionality classes and evaluation of physical random number generators
    - BSI and NIST have different validation processes.
        - Definitions, requirements, modeling, and evaluation processes are different
        - Jointly working to align the RBG standards
        - Will publish a joint NIST-BSI report to explain the process

Part 2: Entropy Estimation

- What is meant by entropy?
    - Commonly we use "entropy" to mean:
        - A string of unpredictable bits
        - A measure of how unpredictable the bits are
    - We measure unpredictability using min-entropy
        - Consider the most powerful attacker you can imagine who might care about trying to attack your system. Consider the probability that they're able to guess the next output from your entropy source. That's what we mean by min entropy. The probability that the attacker can get the next bit, byte, or string.
        - We're trying to bound this maximum probability (P_{MAX}) of guessing the next output, given all possible attacker knowledge
    - We get entropy from an entropy source
- Big picture: Entropy and SP 800-90
    - Deterministic RBGs
        - These are built using a deterministic algorithm
        - The algorithm is published so the attacker knows everything
        - It is based on a cryptographic primitive
            - The ones we have are based on either block ciphers or hash functions
            - Could be based on other things
        - Takes an unguessable string as input and produces a string of indistinguishable-from-random output bits
            - If the attacker can't guess the input string, then the output will be impossible for the attacker to distinguish from perfect random bits
    - We need the entropy source to provide:
        - A string of bits
        - A known amount of entropy
        - Internal tests to make sure it's working
    - Everything in the system, except the entropy source, is deterministic.
- How to build an entropy source
    - We start with a noise source. This is where entropy comes from.
        - Provides us with bits that have some unpredictability
        - Our job in entropy estimation is to figure out how much unpredictability is coming out of this source
    - We perform health tests to verify the noise source is working correctly
        - This is important especially if you're doing conditioning on the bits.
- Two types of noise source
    - Physical source
        - Intentionally built to provide entropy

        - Its unpredictability is generally based on a physical phenomenon that is well understood
        - The entropy source should be simple enough that you can model it and you understand that model.
        - You need to be able to quantify how much unpredictability, how much entropy, you're getting from it
        - Real world examples: ring oscillators, metastable latches, noisy diodes, single photon sources
    - Nonphysical source
        - This is a found source of entropy. You haven't built it to be unpredictable. You've observed that it is unpredictable based on some physical phenomenon
        - Typically measured on a computer in software
        - Examples are interrupt timings, memory access timings, hard drive access timing
            - Hard drive access timing does not work with solid-state drives
- Estimating entropy
    - To evaluate them we need a lower bound on min-entropy and an upper bound on the probability of a very powerful attacker.
    - We can estimate entropy in two ways:
        - Modeling
        - Statistical testing / black box estimators
    - Both are needed
        - We currently lean more on black box estimators
        - We are trying to move more to the modeling
            - Provides a better result but is harder to do
- Statistical testing / black box estimation
    - Requires raw bits from noise source
        - It's not always easy to define exactly what raw is so you need to collect a lot of data
    - Then you need to figure out what kind of entropy estimate you want based on whether your source is independent and identically distributed (iid) or non-iid.
    - This works better as a sanity check rather than as a direct entropy estimate
- IID sources
    - If the source is really well behaved, then every output is independent of all other outputs not varying over time
    - This makes entropy estimation very easy
        - Just count the most common output and apply a binomial bound
    - Most sources are not iid
- IID evaluation
    - We only consider a source as possibly iid if the designer claims it is iid
    - We can use a complicated set of tests to try to falsify a claim of iid
    - An iid claim must also be justified in the report. In other words, the reviewer must verify this is a reasonable claim for this source.
    - If accepted as an iid source, entropy estimation is simple
- Non-iid sources
    - Most sources are not iid
    - Even if the source passes iid tests, it may not be reasonable to assume independence of nearby outputs
- Black box testing

- For non iid sources, it's kind of an ad hoc process. Similar to throwing spaghetti at the wall and seeing what sticks
- We have a large set of these black boxes. Each one is an algorithm that examines a sequence of outputs from a source. Based on some internal set of parameters and some models that it starts with; it estimates what the entropy is.
  - The best estimators are the ones that are based on predictors
  - We take the lowest estimate from all the black box estimators
- We collect sequential data and restart data
- We derive estimates from both sequential and restart data
- Results are probably conservative but it's still very ad hoc and hard to justify
- Black box testing without knowing internals of entropy source is not very powerful

- Modeling
  - You start with a complete understanding and description the source
  - Then you build the stochastic model
    - This is the best process
    - We're working with the German government on their standards
    - You build a model to describe source behavior
    - Estimate the parameters of the model
    - Drive upper bound on maximum probability from the model and lower bound on minimum entropy
    - This is only practical for physical sources
  - Evaluating this is a lot of work. It takes some expertise and that's been a problem for us. We are working on solving this problem
  - There is less rigorous justification for non-physical sources
    - Use a more heuristic approach
    - You describe measured behavior and experiments and try to justify the existence of the entropy
    - Ultimately there's much less of a solid sense since you can't really trace this to some physical phenomenon
- Questions about the noise source
  - This is a set of questions to act as a starting point for thinking about modeling noise sources
  - The first question is just how does the noise source actually work?
  - The next question is where does the unpredictability come from?
  - How much entropy or output is produced?
  - How do you know? You need to justify the entropy estimate.
- Estimating entropy: summary
  - We need to know how much entropy we're getting from the noise source
  - There are two ways to do this: modeling and statistical testing / black box estimation
  - Questions to start with:
    - Where is the unpredictability coming from?
    - Can I quantify it?
  - The black box estimators are a sanity check but can be badly wrong
  - Model estimate is better assuming your model describes the source well

Part 3: Validation

- Validation process
  - To comply with FIPS 140-3, all DRBG's and entropy sources must be validated

- DRBG's are validated through the cryptographic algorithm validation program (CAVP)
  - This is a fully automated validation using tests that we host on our servers
  - We publish a public validation certificate
- Entropy sources and RBG constructions are validated through the cryptographic module validation program (CMVP)

- Entropy source validation
  - Some history:
    - Beginning on November 7, 2020, entropy sources in FIPS 140-2 and 140-3 module submissions are required to be compliant with NIST SP800-90B
    - Until mid-2022, all entropy reports were submitted along with the module validation report
    - Beginning mid-2022, entropy source validation test system (ESVTS) is available for separate entropy source report submissions
- Current entropy source review process with module (until January 1, 2023)
  - Shows an example using a module vendor that builds, designs, and sells cryptographic modules
  - They contract with an accredited testing laboratory to test their module which can be software, hardware, or a combination package
  - The same lab is used for both module and entropy report
    - In the future, there will be a separate cost for entropy with dedicated entropy reviewers
  - The test report is submitted to the C MVP
    - The report goes through review-comment cycles
    - Once finalized, the module with the entropy source is validated and assigned a certificate number
- Entropy source review process after January 1, 2023.
  - We are decoupling entropy source validation from module validation
  - The lab submits raw noise, restart samples, and any conditioned output sample data through E SVTS
    - 800-90B RGB estimators run on ESV servers
  - Once the review is complete, we issue the certificate which can then be referenced in a similar manner to the CAVP certs
  - They can then reuse validated entropy sources in multiple modules. This helps prevent duplicate reviews
  - To support this, we have a web-based client, a programmatic interface or protocol for submitting reports, as well as a user interface where labs can enter of the data for this process
    - All this information is publicly available
- Observations
  - Entropy source validation is a review intensive process
  - Entropy report reviewers and lab staff need specialized technical backgrounds
    - Digital and analog circuits, semiconductor physics, information theory, stochastic processes
  - There have been 44 entropy sources validated up to September
- Conclusion
  - Generating random unpredictable numbers is hard
  - Many things can go wrong either intentionally or unintentionally
  - Standards and guidelines are useful, but they have limitations
  - A good understanding of the design is necessary to estimate entropy
  - Development of guidelines on random number generation is an ongoing process
  - Contact information: rbg_comments@nist.gov

Discussion

Question from Ms. Moussouris on guidance on revoking certificates

- She asked if there is guidance or even mechanics for revoking the certification if research reveals a different way of testing that downgrades the quality of the entropy itself. Are you providing guidance on implementations that would facilitate algorithmic agility in the source of entropy?
- Mr. Hall replied that they do have a well-established process for revoking certificates. Relating to the idea of crypto agility, I have seen some modules that use multiple entropy sources. Perhaps that gets to your concern.

Question from Ms. Fanti on the use of machine learning and deep learning for testing entropy sources

- She asked about the use of machine learning and deep learning for black box testing of predictive models. She also asked, for cases where they don't have a physics-based model, is there any utility to using deep generative models as a computational proxy for those physics-based models.
- Mr. Kelsey replied that one of his coworkers, Carrie McKay, has done a fair bit of work. Probably the best general estimator we have is a high depth Markov model. Carrie has worked on machine learning algorithms, trying to find a way to get better predictors with no success so far. We'll be looking at that and seeing if we can have more general ones.
- Miss Fonte clarified her question. She asked when you derive physics-based models, or when the entropy source provider creates physics-based models, are they deriving closed form expressions for the upper bound on P Max? Are they computational?
- Mr Kelsey replied that it depends on the specific device. Sometimes people can write down a closed form model for this approximation. He gave an example that if you have a source that is based on how many clicks the Geiger counter records for a period, that's something where there's a fairly nice twist on distributions. You then might have to deal with the fact that reality doesn't ever perfectly fit your model. You may have to do some simulation or numerical calculation to get the estimate at the end. It's still useful because you're able to track your assumptions back to the device and the parameters of the device.

OMB M-22-18 Enhancing the Security of the Software Supply Chain through Secure Software Development Practices

Mitch Herckis, OMB

- OMB M-22-18 related to EO 14028, Improving the Nation's Cybersecurity, which was released May 12, 2021.
- This is part of a much larger effort underway within the federal government around securing our broader nation's cybersecurity
- Is an area that both the public and private sector is trying to figure out how to do well, balancing the need to reduce the risk to all enterprises while still delivering services

Background

- There has been and continues to be a lot of focus on supply chain security. It continues to be a major threat factor.
- EO 14028 laid out a clear and urgent message that we need to reduce the risk from our adversaries to potentially compromise US systems.
  - Section 4 addresses secure software development

- We are talking about section 4(k) which is a small piece of a very broad implementation plan
  - Part of the plan was for NIST to put forward secure software development standards
    - Those were released as SP800-218 and the related NIST software supply chain security guidance which explained how federal agencies should implement 800-218.
  - 4(k) has requirements for OMB to "take appropriate steps to require that agencies comply with the secure software development standards that were outlined by NIST".
  - OMB announced compliance requirements in March, but we recognized there was a very complex issue of how to deal with vendor attestation.
    - It is understood that there are software producers who create most of the products that our government runs on.
    - We worked with our partners at NIST and saw broad input from the private sector as well as from agencies.
    - We held a public workshop; we collected written responses to a key set of questions that we saw as outstanding and critical to success.
    - We got lots of quality responses
    - The result of this effort is M-22-18

What does it do?

- It directs agencies to use third party software that complies with minimum standards of the Secure Software Development Framework
- It provides for the creation of a common self-attestation form for software producers and agencies
  - Hoping this will reduce burden and red tape
  - Will allow the federal government overtime to quickly identify security gaps when new vulnerabilities are discovered

Timeline

- Applies to agencies' use of certain third-party software that's developed or has had major version changes following September 14, 2022.
- On June 11th, 2023, agencies need to start collecting attestations for critical software
  - Critical software is defined in a separate OMB memorandum
- On September 24, 2023, agencies need to collect attestations for all other software that's covered

What's covered and what isn't covered?

- Starting with what isn't covered and does not need an attestation
  - Software developed in-house with your own staff
  - If an agency is using an open-source software library or free tools
    - We recognize that, with open source, if there's not a relationship of some sort between the software producer and the agency, it could become very difficult to track down and have someone submit an attestation
    - This does not mean we are not paying attention to that issue, but it is not covered here
  - If an agency has used software that was developed prior to the release of the memo and hasn't undergone any major version changes, there's no attestation necessary
  - If an agency is using software acquired or licensed under another agency's contract and the other agency has already received that attestation
- What is covered?
  - If an agency uses software, whether it's self-hosted, software as a service, or some other cloud hosted service, that was developed after the release of the memorandum

  - If an agency begins using a new major version of the distributed software that was released after the issuance of the memorandum
  - If there is a software as a service product which has continuous updates, we are building in a process for that where the vendor will need to attest to conformance and will be responsible for notifying the agency
  - An agency that uses a software bundle where one or more pieces of that bundle were developed or had major version changes after the release of the memorandum, that would apply for those pieces, and they would have to attest for those pieces.
- There are extensions and waivers allowed.
  - Granted at the OMB director level
  - There is a high bar
  - Waivers are granted in exceptional circumstances and for a limited time
  - The agency must also submit a plan for mitigating any potential risks along with the request for waiver
  - In addition, the agency can request an extension for complying with specific requirements under the minimum standards.
  - Request needs to be transmitted at least three days before any relevant deadlines and associated with a plan of action and milestones comment explaining how the agency intends to meet the requirements

Additional Guidance

- The Federal Acquisitions Regulation (FAR) Council will be issuing some guidance to agencies on how to work these things into their acquisition process.
- Right now, the self-attestation form will be issued by CISA.
  - The form will list out specific secure development practices and tasks that OMB is responsible for identifying and
  - Software producers will be required to attest if their development process incorporates those elements.
  - These are these are minimum elements.
  - Agencies have the right to say because we value this piece of software and we recognize the specific risks to this environment, we are going to require something additional.
    - Could be an artifact like a software bill of materials or digital pieces of the Secure Software Development Framework.
  - If the software producer has business across multiple agencies and sub-agencies, they can reutilize that form.
  - Agencies can consider reciprocity of the form.
- Software producers are allowed to use the standard self-attestation form or to submit a third-party assessment provided by either:
  - a certified FedRAMP third party assessor organization or
  - one approved by the agency in lieu of a self-attestation.
  - Allows for some flexibility there.
  - Example: Perhaps you use a lot of open-source software, and you want to go the extra length and say look at the way we package this together. They want to show, through a third party, that they're doing the right things.

Government-wide repository

- The memorandum says CISA will develop a government wide repository for attestations and other secure development artifacts, such as software Bill of Materials.
- Our goal is to allow for common usage search, storage, and security measures for the software developer.
    - Can be stored in a secure manner and be searchable so if agencies or CISA are looking to rapidly determine what in an environment might be of concern, there's a central way to handle and manage that over the long term.
    - That won't be immediately available, but we want to ensure we are on the path to building as soon as possible

Question from Dr. Baker about Third-Party Assessments

- Dr. Baker asked if they will be requiring third-party assessments from an independent organization.
- Mr. Herckis replied that they are not requiring that for a few different reasons:
    - One is they don't believe it's necessary in all cases,
    - Two, they want to reduce the burden on agencies out of the gate to ensure they are able to implement this as soon as possible.
    - And they're uncertain about the timelines that will create.

  He added that it's uncertain what additional value that might bring. There could be use cases where that would help but remember, they're submitting to the government a form signed by senior professionals within their organization attesting to certain standards. It is something that hopefully software producers have taken relatively seriously before signing on that bottom line.

- The Chair added that there's probably a longer debate. From his experiences, software security comes from the actions of the vendor that's building and protecting the software. There are two possible outcomes with a third-party assessment, one of them is that you get an in-depth review, and you get higher assurance. Another possible outcome is that you get a document, basically outsourcing your assurance to somebody who doesn't understand what you've done. Because of the detail and reliance on internal processes and software security, he thinks that's the more likely outcome of third-party assessment.
- Mr. Herckis added that agencies have the right to go beyond the minimum requirements.

Question from Ms. Moussouris about attestations and Open-Source Software (OSS)

- She asked how far down the dependency tree do you go for attestations?
- Mr. Herckis replied that software producers have to attest to anything that's not open source.
- She replied that her second question was why is OSS exempt?
- Mr. Herckis replied he doesn't have an answer for that.
- The Chair added that OSS is covered under the EO.
- There was a general agreement that OSS is covered somewhere in the EO but no consensus on where.
- Ms. Moussouris added that she is concerned that not including open source could inadvertently incentivize agencies to choose open-source solutions because they do not require attestation.
- Mr. Herckis indicated that someone asked him if they would be able to use open-source solutions ever again. He thinks agencies are going to want to use what they need for the scenario to get the job done. We are trying to encourage this to be a pathway that is easy and there is a way for this to be something that both the software producer and the agency can utilize. The underlying goals here is to reduce that risk, while not interrupting the business processes that are necessary to continue agencies delivering services for the American people.

We know it doesn't do all things and it some people have told us it doesn't go far enough. Others have told us this is going to be burdensome and goes too far.
With open source, if you don't have an actual person who's in charge of it, it's much harder to get some sort of attestation.

Comment from Mr. Gattoni on version numbering and using the numbering to avoid having to submit an attestation

- Mr. Gattoni indicated that he worries that the import of a bug fix for a very critical vulnerability could fly under the radar if the version change fails to align with the significance of the update. He recommends zeroing in on the specificity of why a version change is numbered a certain way, and some credibility with the software developer, whether they're doing that for the build of the software or marketing reasons.
- Mr. Herckis agreed and indicated it is a focus because of the scope of how things are built when it comes to software development, and then various iterations and different types of elements. We're trying to ensure that we're taking the right steps with the right pieces and building along that path and maturing as we go.

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

Public Comment, Summary of Day 1, and Board Discussions

- No public comments were received.
- Board Discussion
- AI Risk Management Framework
  - Mr. Groman comments:
    - Needs more concrete guidance to be a useful document
    - Need to be specific when saying things like “needs to address impacts on civil rights.”
      - What does that mean. Give examples of potential negative outcomes.
      - It's not yet implementable.
      - Need a more granular, practical set of potential negative outcomes to consider.
    - Bias is a very broad term. Bias in and of itself isn't negative or positive.
      - Again, needs to be more grounded in context with examples.
    - There should be no subjective judgement in the document itself. The document itself doesn't determine what's fair, and the document itself won't assess whether certain negative outcomes in a given situation outweigh the positives. Those subjective judgments are made by the organization.
    - What's missing is the factors in a practical way to consider so the entities in different contexts think through the potential negative impacts.
    - Would like to see more cross-pollination between the frameworks where they all have a catalogue of adverse outcomes similar to what is in the PF.
  - Mr. Venables comments:
    - Also missing that layer of who's supposed to do what with what? Activities and responsibilities assigned by role.
      - For example, on this AI point that you bring up, you can imagine if there's a set of guidance specifically for the engineers and data scientists and others that are implementing things that has to be suitable for them. Then there's a broader framework in which other things need to be addressed, but they can't be addressed by the engineers.

They're addressed by some kind of risk governance AI ethics committee of the organization.
    - It'd be a useful construct going forward for all the frameworks to just have a little bit of an overlay indicating which bit of which section, which goals, should be typically undertaken by which role in the organizations.
  - Ms. Moussouris mentioned reaching out to a company she knows who is doing something like what Mr. Venables described. She offered to reach out to them for more information.
  - Ms. Flynn Goodwin comments:
    - It would be interesting to find out how the AI RMF is being socialized with small and medium-sized businesses. These businesses don't have the same level of resources. How would they implement these things?
    - The Chair indicated that it's possible that smaller businesses may not be tracking what NIST is doing.
    - Ms. Moussouris added that even small businesses may have millions of customers and thus have the data protection burden and need to be held accountable.
    - Ms. Flynn Goodwin added that her concern is how to raise awareness of this with small companies that don't track this space?
    - Action Item: The Chair asked Mr. Scholl if he could communicate these concerns back to NIST. Mr. Scholl agreed to bring these issues to the attention of Mr. Romine and others at NIST.
  - Mr. Venables would like to do a follow-up on the notion that the frameworks need to be mapped to the roles of people that consume them. This could be a letter or a topic for future meetings.
- RMF, CSF, and PF
  - The Chair's comments:
    - Need to follow up on this for more discussion.
      - It's too easy to get confused on how they relate to each other and how to use them together.
      - There's also the related matter of trying to make the control implementations as consumable as possible, which is something we didn't hear about today.
        - A combination of secure configuration and secure checklists, and useable guidance
- OMB M-22-28
  - The Chair indicated that this is a topic he would like to have more conversation on.
  - It was agreed that they would continue talking with OMB on this topic.

Day Review and Meeting Recessed

The Chair adjourned the meeting at 4:30 P.M. ET.

Thursday, October 27, 2022

The Chair opened the meeting at 10:01 A.M. ET and welcomed everyone to the call.

NIST Cybersecurity and Privacy Update

Matthew Scholl, NIST

Kevin Stine, NIST

Frameworks

- Mr. Stine expressed excitement with how the frameworks are helping advance the discussion in different places
- Each framework tries to focus on different sets of challenges in managing risks, although there are similarities between them as well.
  - Need to highlight those similarities and points of intersection as well as the differences.
- There was a discussion at the end of the day yesterday on the degree to which the different teams interact.
  - There is much cross-pollination. A lot of the staff on the CSF are also writers on the AI RMF.
  - The same can be said for the CSF and PF.
  - There are also staff commonalities and technical experts from the secure software development framework efforts on the other frameworks as well.
  - We need to do a better job of communicating that cross-pollination, how they relate and how they're different.
- CSF Update Process
  - Doing lots of engagement through conferences and roundtables with government and industry stakeholders.
  - Includes interactions with foreign governments as well including Italy, Colombia, and Singapore.
  - Organizing the second workshop for the early part of next year. No specific dates yet.
  - Hosted a virtual workshop in September.
    - Very interactive with good attendance
    - Great source of input
    - Had participants from about 100 different countries
  - Future meetings will be public, probably hybrid in-person and virtual.
  - Will continue to seek public comments on drafts.
  - Continue to see greater uptake and applications of the current CSF
    - Recently issued a CSF profile for liquefied natural gas.
      - Collaboration with the Department of Energy (DoE)
      - Helping identify priorities, cybersecurity outcomes, vendors with standards and guidance in the context of liquefied natural gas business processes.
    - Recently launched a new project with DoE on developing a CSF profile for extreme fast charging stations for electric vehicles.
  - Mr. Scholl reiterated that there is a great deal of cross-collaboration on CSF 2.0. He's working on the breakout for updating the concept of the tiers in the framework.
- Cybersecurity metrics and measures
  - The CSF is one of the ways they are looking into cybersecurity metrics and measures up and down the abstract stack.

- Referring to the entropy discussion yesterday, at the low end of the stack, we abstract a lot of our security on something as foundational as encryption.
  - We ask, "are you encrypting, or do you have a good tunnel?". The assumption is that the encryption is good.
  - The assumption on top of the encryption is that your entropy is good.
  - The assumption on top of the entropy is that your noise source is good.
  - As a result, sometimes we are building these abstractions that we use for our measures and metrics on foundations that we are still struggling to measure.
- Looking at programmatic measures down through specific things like forensics work NIST does in areas like the:
  - National Software Reference Library
    - Take individual pieces of software and hash them for people to measure and understand where that software is.
  - National Vulnerability Database
    - Assign metrics to vulnerabilities
  - Want to establish some common terms and understandings of everything we've been talking about.
    - Plan on having a workshop this winter and then publish some common understandings around what we mean by "measure" and "metric".
      - Ordinal metric, qualitative metric, comparative metric
  - Looking at all the ways other organizations are implementing these things, such as a CSF implementation

- **Mr. Venables** asked if there is a formal process for prioritizing future work. Part of the impetus for this question is:
  - We're seeing a lack of research in operating system security and the evolution of hardware, specifically embedding security features in hardware.
  - There's a lot of pressure on the operating system to make up for security deficiencies in hardware

  Is there a role for NIST to, at least, set a directional framework for the future of hardware security and what the OSs need to do?
  - Mr. Scholl replied that there are a couple of formal processes:
    - We get R&D priorities from OSTP and NITRD.
    - He is a co-chair of the Cybersecurity Information Assurance Working Group within NITRD.
      - Put together an annual cybersecurity R&D plan for the federal government
    - Other drivers are legislation and Executive Orders
    - Focus on nine general areas that establish the guide rails
      - Perform bottom-up reviews
    - They are very interested in hardware and hardware security (see CHIPS Act)
    - Have not been as focused on OSs but they have been doing a lot of work on containers and virtual environments
    - They have a significant amount of work in software being driven by the EO.
    - Have not done a lot of work to stitch all of this together.
    - Mr. Venables agreed that the work around interweaving all of this together is important and suggested this might be worth a workshop.
    - The Chair agreed with Mr. Venables
- **Ms. Moussouris** asked if they are able to measure which industries are lagging in implementation, especially if they are part of critical infrastructure, and why? Do they need more resources?

Guidance? Could we get a "heat map" on who is hitting the lowest maturity level in the implementation tiers?

- Mr. Scholl replied they don't measure that across industries, but they do rely on the sector agencies to let them know what the sector needs. Also lean on the National Risk Management Center to identify critical needs.
  Another internal debate they have is about the size of an organization and how that affects the challenges they face. Maturity may be more of a factor than size.
- Mr. Stine added that he believes that type of data would likely come from an organization like CISA, probably with the sector risk management agencies.
  He agreed that they would benefit from this kind of information to help inform improvements or identify gaps.
  A lot of feedback on the CSF revolved around usability.
  Turning to the community for help identifying opportunities for measurement as well as maturity. Maturity could be within an organization, between organizations within a sector, or possibly between sectors. The don't have answers yet but agree it needs to be part of the discussion.
- The Chair added that there seems to be a gap or disconnect between the PF, CSF, and RMF. The RMF is what's mandatory for a lot of agencies and organizations. It's important for them to be more integrated.
- Mr. Scholl agreed that they can do better and provide more clarity, especially since the CSF is also required now.
  - Need to provide better tools
  - Have a meeting scheduled for December 6 with a couple of agencies that will share lessons learned and how they are working both the CSF and RMF. If that meeting is successful, maybe they can talk at an ISPAB meeting.

- **Ms. Flynn Goodwin** mentioned the EU Cybersecurity Certification Framework that will impact a lot of U.S. companies. She was wondering if they are also assessing progress there and if there are risks/benefits to harmonizing?
  - Mr. Stine replied that they are tracking that. It's a different part of the Department of Commerce, looking at helping U.S. businesses be competitive around the world.
  - **Ms. Flynn Goodwin** feels this may be an opportunity for exploration. There has been interesting dialogue with Germany and France pushing hard for certification schemes.
  - Mr. Stine added that they are seeing relations between this and other technology domains like IoT.
  - Mr. Scholl brought up how they're working with standards development organizations such as BSI Germany and ANSSI at ISO. For non-technical issues (policy, trade, etc.) they work through the Dept. of Commerce and the International Trade Administration (ITA) and the State Department.
- Crypto
  - Will get an update later on the post quantum work
  - Continuing research into new privacy enhancing types of encryptions
  - Will be looking at anonymous and blind signature capabilities
  - Continuing research into partial and fully homomorphic capabilities
  - Interested in zero knowledge proofs and threshold cryptography
  - Over the next 2-10 years it's going to be a turbulent time for encryption
    - Coming up with a timeline for removing long-standing cryptography and adding new
    - Will be deprecating things as well.
      - Will be dropping 3-key Triple DES

    - Looking at SHA-1 and RSA
    - Getting into the standards, protocols, products, and structures
  - **Mr. Venables** mentioned the crypto-agility paper they produced and emphasized this is an area to continue to bring into our standards. Crypto-agility needs to be a core of what organizations are doing.
  - Mr. Scholl replied that they will be updating some of their core documents to remove ambiguity.
  - **Ms. Moussouris** mentioned that as OMB is enhancing the security of the software supply chain through secure software development practices, they should also be pushing crypto-agility as that will dictate what the government adopts.
  - **Ms. Miller** reminded Mr. Scholl that defense has large weapons systems and embedded systems to keep in mind.

SP 800-63 Update and Plans for Identity

David Temoshok, NIST
Ryan Galluzzo, NIST
Connie LaSalle, NIST

What is the state of digital identity today?

- Pandemic drove rapid transition of services from analog and in-person to digital
- New vectors and opportunities for nefarious actors
  - Allowed individuals to get access to benefits they needed but also opened the door to new attack vectors and opportunities
- Advanced and scaled threats – leveraging ransomware, building out synthetic identities, developing automated attacks to go after registration systems, black market and dark web areas with lots of information available
- Many systems are being protected by outdated techniques.
  - Passwords and knowledge-based verification
- All this has led to some very negative fraud outcomes
- Limited alternatives – there are some applications and new technologies to increase the security and assurance of online transactions, but they are often very tightly focused technologies or ones that have some controversy or issues, such as facial recognition, or resulted in long identity proofing processes.
- Changing consumer and public sentiment
  - individuals are becoming more aware of:
    - their online presence
    - Risks and threats to personal privacy and safety
  - Also want to have a better understanding of what's happening with their data once they start to engage in some of these transactions.
- Not seeing the interoperability anticipated
  - Federation is still not the ubiquitous thing that we were hoping to see at this point.

NIST's Role in Identity Management

- Focused on NIST's ongoing identity projects:
  - Updating NIST SP 800-63, Digital Identity Guidelines
    - Brought in folks from privacy, usability, and biometrics teams to ensure they address all risks, not just security risks
  - Updating NIST SP 800-157, Guidelines for Derived Personal Identity Verification Credentials

        - Want to provide more options, flexibility, and security to federal agencies as they expand beyond traditional PIV implementations.
    - Creating new guidelines for PIV Federation to promote greater cross agency interoperability
    - Developing mobile driver's license guidance in close collaboration with DHS TSA, ISO, and IEC.
    - Researching identity verification and attribute validation technologies to promote better understanding of the technologies and how they can be applied to our identity proofing guidance.
        - Also looking at verifiable credentials and other methods that might be used to help provide strong assurance
    - Developing Zero Trust (ZT) reference implementations to advance critical national cybersecurity priorities.
        - NCCoE recently updated their ZT guidance

What are the Digital Identity Guidelines?

- Initially published in 2004. Paralleled the government's movement of putting services online so the public could access services online.
- Revised three times. Most recent revision was in 2017.
- Presented the processes and technical specifications for online access:
- Current revision:
    - Bring current with technology and against attack vectors
    - Take into consideration lessons learned from implementations across the federal government and outside organizations
- Four volumes:
    - First volume deals with applying risk management principles to digital identity management
        - As part of the RMF, 800-63 doesn't specify a single level of controls.
            - Recognizes that the assessment of risks and determination of appropriate controls are graduated.
            - Apply the risk management approach for establishing baseline levels of controls to mitigate low, moderate, and high risk across a broad range of functions
            - Baseline controls are called assurance levels
    - Volume 63A addresses identity proofing and enrollment into an identity system
    - Volume 63B addresses authentication: how they can prove their identity to access online services without being ID proofed again.
        - Addresses how persistent authentication and access to the government's online services can be accomplished for the public.
    - Volume 63C deals with Federation
        - The sharing of authentication information and personal information from one agency or security domain to another.
        - Allows the sharing of authentication status and information across government in a secure and privacy enhancing way
        - Encourages Federation as a shared capability across the federal government.
- Implements FISMA and, thus, are mandatory to be implemented across the federal government
    - May be voluntarily adopted and implemented by organizations and industry outside the federal government
- **Question from Ms. Flynn Goodwin**

- She asked what their view is on the ability to use data about the state of identity across the federal enterprise to help inform best practices of what's currently being deployed and where we need to be as a model of digital transformation? Do they have data from each of the federal agencies that reflects where they are implementing this standard and to measure its effectiveness? Something to put into the public domain to give technical credibility to the implementation mandate.
- Ms. LaSalle replied that in revision 4, they have taken into consideration the challenges and opportunities faced by federal agencies in the last 5 years to include their service providers and contractors and to have a document that reflects the real world.
- Mr. Galluzzo added that they don't have a specific set of 800-63 reporting metrics or requirements. Some aspects are captured in the FISMA reporting. They have spent a lot of time talking to the agencies but that is a good point.

Why Are We Making Changes?

- 2020 issued a call for comments on changes and experiences in implementing the 2017 version.
- Discussions with agencies and industry on implementations, lessons learned, and their experiences.
- Needed to address equitable access for underserved communities to government online services per EO 13985
- Provide alternatives and optionality for choice for agencies, organizations, and individuals
- Ensure currency with controls to deter fraud attacks
- Ensure the guidelines address real world implementations based on feedback from agencies and industry
- Expand risk assessments beyond identity teams to include interdisciplinary teams, considering the population being served and the need for equitable access
- Clarify and consolidate requirements

What We Aren't Changing.

- Publication structure – still 4 volumes
- De-coupled assurance levels – still 3 types of assurance levels with three levels of assurance
  - IAL – Identity assurance level
  - AAL – Authentication assurance level
  - FAL – Federation assurance level
- Assurance levels are still determined by looking at the risks associated with each of those three areas of digital identity and then tailoring the controls to the agency specific risks

What We Are Changing.

- Adding equity considerations
  - Each volume has separate sections dealing with privacy, usability, and equity considerations for the processes and controls to make sure that implementation decisions reflect this full range of considerations.
- Base volume changes
  - Feedback indicated that rev 3 was too compliance focused
  - In the new version, they tried to be more process oriented and align better with the RMF
    - Focus on understanding the risks and using that understanding to determine the base assurance level
    - Integrating security, privacy, equity and specific threat considerations all into the equation of tailored assurance levels for identity proofing, authentication, and federation

- Created non-federated and federated overarching models of the process
- Added information on integrating multiple viewpoints in the risk management process and thinking about how to establish an end-to-end continuous management and continuous improvement process.
  - Identity should be integrated with response, cybersecurity, pricing, and other teams so if there is a failure of identity controls that results in something happening that shouldn't, there is a way to tie it back in an create updates to your overall program.
  - Can add mitigating and compensating controls to back-end process to free up some friction on the front-end to support greater access and usability.
- Included considerations to ensure that social impacts to individuals and communities are considered when making decisions around technology and the implementation of technology.

- Volume 63A changes:
  - Changed IAL1 to be able to address low risk access that still requires some confidence in the identity of the individual accessing the online service.
    - Parallels the processes of the higher IALs, just with less rigor to allow greater optionality for agencies and individuals
  - Trusted referees and applicant references:
    - Intended to provide capabilities to individuals who otherwise would be disadvantaged in the ID Proofing process
      - Can provide support to individuals at the process level or if they have difficulty meeting the documentation requirements
    - Trusted referees are agents of the federal government or the ID service
      - Specially trained to be able to make risk-based decisions based on unique circumstances for the individual that may have precluded them from meeting the process or evidentiary requirements at any of the IALs.
    - Applicant References deal with individuals who may not have the capability to complete the process or evidentiary requirements.
      - Can vouch for the individual
      - Are familiar with the circumstances faced by the individual
      - Able to vouch for information and the unique circumstances of those individuals to allow trusted referees to make risk-based decisions to allow them to participate in the government online services.
    - Adds considerations for digital evidence
    - Updating biometric performance requirements
      - Providing for additional optionality beyond comparison of a facial image to a facial portrait on ID evidence to allow the binding of that evidence to the applicant and allow greater flexibility
  - Adds "Subscriber Account"
    - As information is collected, it's important for individuals to be able to access that information and know what information is being collected.
- Volume 63B changes
  - Added detailed description of phishing resistance
    - **Mr. Venables** interjected that there are organizations deploying stronger forms of phishing resistant authentication, like FIDO keys, and attackers are shifting to post-authentication attacks using cookie theft. Industry response is to introduce bound post-authentication credentials. Is that kind of thing in scope for 63B?

        - Mr. Galluzzo replied that 63B covers the authentication component but does include session management requirements for controls that address ongoing session protection.
        - **Mr. Venables** asked if NIST is doing anything with W3C or IETF? There are changes needed in other standards and it would be useful if NIST could help drive that change.
        - Mr. Galluzzo and Mr. Temoshok replied that we have team members participating in IETF and ISO working groups and, once we get the draft 800-63-4 out, one goal is to engage better and make sure what's in here ends up in technical specifications, international standards, etc.
    - Password management
        - **Ms. Moussouris** asked about password rotation, indicating that this can lead to password reuse with small changes. She also asked about password lengths.
        - **Mr. Galluzzo** replied that they have included things like "should not have password rotations" for the exact reason she mentioned.
        - They do specify a minimum length, with the focus on creating long passwords rather than more complex ones, in conjunction with the use of block listing which restricts the types of words and characters that can be used as well as the use of a multi-factor scheme.
        - **Ms. Moussouris** mentioned that a good resource for data breach information is the Verizon breach report.
- Volume 63C updates:
    - Federation is sharing authentication status and information from one security domain to another. Individuals don't need to be ID proofed everywhere they go
    - One aspect of Federation is the identity service (the identity provider or IdP) is disclosing authentication status and personal information to the relying party (RP) therefore there must be an agreement between the RP and the IdP on how that information is disclosed and what's done with it.
        - Update includes process requirements, controls, and guidance on how they should be implemented and presented in the federation trust agreement. This forms the core for the federation services.

Discussion

Question from Ms. Fanti on biometric performance requirements

- She asked how these requirements account for differences in performance. E.g., different types of skin tones
- Ms. LaSalle replied they're focused on discrepancies in performance rates, ensuring that performances are equivalent across demographic groups.
- They have created metrics that are specific to identity proofing and verification.

The Chair recessed the meeting for a 10-minute break.

Cyber Safety Review Board Log4J Report

Abby Deift, HQ DHS

Katie Moussouris, Luta Security

What is the CSRB?

- Established through President Biden's EO 14028 on improving the nation's cybersecurity to review significant cybersecurity events so government, industry, and the broader cybersecurity community could better protect against the nation's infrastructure.
- Not a new concept

  - When there is a significant cyber event there is usually an after-action review to determine what happened, where the controls failed, was there anything that could have been done, how did the response go, effectiveness, etc.
    - Generally, that kind of report isn't shared outside of the individual company and organization other than with lawyers, regulators, auditors, and that kind of thing
  - Since cyber events generally don't only impact one organization, there's a lost opportunity for other organizations to benefit from that kind of "lessons learned" and they may also have been impacted by the same vulnerability, event, or issue.
  - The board was meant to create an opportunity to come together to develop a single authoritative set of facts, findings analysis that can be shared across industry across organizations in a way that meaningfully drives enhancements across different stakeholder groups.
- The board serves as an advisory committee.
- Administered by the Department of Homeland Security
- Composed half of private sector leaders in the cybersecurity community and half federal with representation from ONCD from DOD from FBI, DOJ and DHS
- The chair is Rob Silver, the Undersecretary for Policy at DHS
- Convened in February 2022 for its inaugural tasking, a review of the December 2021 Log4j event.

How does the CSRB work?

- Brings together government and industry leaders to identify lessons learned from significant incidents and deliver strategic, actionable recommendations.
- Tasked with reviews by the President, the Secretary of Homeland Security, or the CISA Director
- The board came up with a set of operating principles in its first meeting
  - Purpose of the reviews
    - Are not about finding blame
    - About getting information, analyzing that information
    - Coming up with lessons learned for the impacted organizations.
- For Log4j:
  - CSRB is engaged with over 80 organizations and individuals.
  - Done on a voluntary basis
  - Board asked people who they thought would have relevant information
  - Issued a request for information
  - All done through the lens of the public interest. It's important that they are able to share the report.
    - Report is available online: https://www.cisa.gov/sites/default/files/publications/CSRB-Report-on-Log4-July-11-2022_508.pdf
  - Challenge was figuring out how to describe the genesis of the vulnerability, the reporting of the vulnerability, disclosure to the researchers and vendors, when and how the first attacks started showing up when it was publicly disclosed, which sectors of the industry in the world were affected. How did these attacks spread around the world?
    - We were not able to get all that information
      - Some sectors were missing
      - Not everyone was able to answer because of legal agreements resulting in them not being able to reveal customers or breaches
  - Resulting report included

        - Recommendations for software manufacturers and others who have to incorporate third party libraries.
        - Ideas on how the government could help
    - Looking at ways to improve future processes and protocols

What did the CSRB find?

- Log4j is an "endemic vulnerability" and unpatched version of Log4j will remain in systems for many years to come
- Many organizations struggled to quickly identify where in their environments they have Log4j.
- Organizations committed significant resources to remediation
- Most interviews mentioned the use of social media for getting effective, real-time updates to enable network defenders
- CISA provided authoritative and valuable guidance to responders through innovative channels, including an online repository of products containing vulnerable code
- The board engaged in open-source software security, looking at the security risks unique to open-source.
- Report includes observations related to training and development for software developers.
- **The Chair** mentioned the initiative at DHS on software bill and materials (SBOMs) that's been going on for a few years. Did you find anything about organizations who use SBOMs doing better?
    - Ms. Moussouris replied that she asked about the use of SBOMs in her interviews and none of them really made use of SBOMs for a variety of reasons.
        - The more mature organizations, couldn't stop to evaluate if their product lines were exploitable; they just upgraded the version of the library just-in-case
        - At the moment, SBOMs were not useful in the rapid response.
        - What made a big difference was the ability to patch quickly and to have the muscle memory of doing the patches.

What does the CSRB recommend?

- Issued 19 specific and actionable recommendations in the following areas:
    - Address the continued risks of Log4j
    - Drive existing best practices for security hygiene
    - Build a better software ecosystem
    - Investments in the future
- Recommendations are broad-ranging, and they are directed to different groups of stakeholders
    - Government
    - Industry
    - Academia
- NIST-related recommendations include:
    - Organizations should invest in capabilities to identify vulnerable systems.
    - Software developers and maintainers should implement secure software development practices.
    - Invest in training software developers in secure software development.
    - Increase investments in open-source security.
    - Examine the efficacy of a cyber safety reporting system (CSRS).
    - Establish a government-coordinated working group to improve identification of software with known vulnerabilities.

What comes next?

- We are hoping to identify where the implementation streams are happening both within DHS and outside DHS.
  - The Secretary is interested in where DHS can lead by example or where there are specific actions assigned to the federal government.
- **Ms. Miller** asked if they have given thought to institutionalizing these recommendations in things like FISMA?
  - Ms. Deift replied that they are in an immature phase of that discussion, but it is something that they are interested in. She asked if there is a specific example or area that they should consider.
  - Ms. Miller mentioned looking holistically across the security posture of systems that fall under FISMA. This plays into becoming tighter on weapons systems and the development acquisition process.
- Want to write a report that drives meaningful change
  - In the tracking phase of how to drive implementation of recommendations
  - Mr. Groman mentioned that several of the recommendations are the same as those in documents from 10 or more years ago. We shouldn't need to repeat the same recommendations over and over.
  - Ms. Moussouris agreed but commented that if they hadn't included those basic recommendations, people would ask why they weren't there. She acknowledged that maybe those recommendations need to use stronger language.
  - Ms. Miller suggested adding a caveat on the recommendations that aren't new.
  - Ms. Moussouris brought up the last ISPAB meeting where CISA did a presentation and federal enterprises had around 1.2 million internet exposed endpoints that had not been patched for known exploited vulnerabilities, some of which are a decade old. She agreed that they need to do something about this but doesn't know if there is a way to say it more strongly.
  - Mr. Groman mentioned that this document really needs to target the cabinet secretaries and decision makers who have the responsibility for allocating resources so that more resources go toward IT security.
  - Ms. Deift replied that the board ended up having in-depth discussions around incentive structures at the individual and organization levels. One of the investments in the future recommendations is for the National Academy of Sciences Cyber Resilience forum to look at incentive structures and adoption of best practices and approaches that our interviewees endorsed as effective and identified gaps in adoption. Another is to establish baseline security measures. More to talk about there and they indicated being open to suggestions.
- **The Chair** asked about the CSRS shown on one of the slides.
  - Ms. Moussouris clarified that it stands for Cyber Safety Reporting System. The proposal is to study the creation of something like the aviation safety reporting system. This is a nascent idea. The board is looking into if this makes sense. There's a description in the report.
- **Mr. Venables** asked how the CSRB decides what's next to review. Are those decisions announced publicly?
  - Ms. Deift replied it's codified in the EO (see beginning of the presentation). The timeline for Log4j was the Secretary made the determination (currently there is no new tasking) and then it is announced. With Log4j, it was announced in early February and the board kicked off the review at the end of February.
- **Ms. Flynn Goodwin** asked what their lessons learned were. What would they do differently from a workflow or investigative process perspective? Do they see the board coming in sooner in the incident response process or remaining post incident reflection time?

  - Ms. Deift replied the board is being very introspective. They are very encouraged by the level of participation by organizations and individuals. They are looking to build out their staff and staff support to help drive evaluation of the information received and the way they are asking and engaging with potential interviewees. Regarding timing, the height of the response to Log4j was around December – January and the board started its review in February which is pretty timely. In general, the board wants to take on timely and current events but it's not an investigative body. It's about fact finding and not placing blame.
- **Mr. Groman** asked about staffing and funding. Ms. Deift reiterated that they are building out the organization to better support the compilation of information.
- **Mr. Groman** asked if they found reluctance from the private sector to participate because commercial entities didn't want to broadcast the fact that they'd had an incident?
  - Ms. Deift replied that they were pleasantly surprised by the level of participation. Some organizations took advantage of information protections that the board offered. There are always limitations to voluntary reporting though.
  - Ms. Moussouris added that a lot of them were willing to respond on their patch and asset management, and all the response types of activities they did. But as far as she knew, they didn't receive one actually telling them that they were breached. As a result, the board has no data on which industries were the least prepared to deal with this.
  - Mr. Venables asked if that might be a sampling issue. Ms. Moussouris indicated there was a broad range of companies, organizations, and federal agencies and they didn't get breach data from any of them.
- Ms. Deift made a closing remark that this report was a proof-of-concept, a lot of the context around some of the information requests, the methodology, the process was the first time doing this. They are in a learning period. They came out of this first review with a good product. Now they are looking at lessons learned and where they go from here.
- Ms. Moussouris mentioned she put her observations up on her blog about what they saw unfolding. One of the things she feels is missing is the board advising the federal government not to adopt the same kind of regulations as they have in China for reasons of not wanting to concentrate so many vulnerabilities into a federal government reporting system of some kind, that would then provide our adversaries with a concentrated target of attack.

The Chair recessed the meeting for a 55-minute lunch break.

Open-Source Security Workshop Readout

Angelos Keromytis, Georgia Institute of Technology
Michael Ogata, NIST

Open-Source Security Workshop (Angelos Keromytis)

- OMB and the National Science Foundation initiative looking at how to improve the state of security for open-source software.
  - Mr. Keromytis was on a steering committee made up of folks from industry, academia, and government in May 2022
  - Brainstormed ideas on the big issues in three areas:
    - Memory-safe programming languages
      - What will drive the adoption of memory safe languages and other similar technologies that could improve the state of protection from the most common types of vulnerabilities.
    - Software dependency management

        - What can we do about the challenges relating to supply chain and software dependency management, e.g., software sprawl?
        - How do we manage vulnerabilities, accidental or purposeful, inserted in through the open-source chain or package?
    - Behavioral & economic incentives to secure the open-source software ecosystem
        - What can we do to incentivize and enable open-source projects and developers to adopt better strategies, better technologies, better workflows, and better tools?
  - Held an online workshop with about 3000 attendees from industry, academia, and a few from government
    - Goal of workshop was to expand on the ideas in the above areas and come up with actionable ideas for the short term and long-term steps that the Government could take to improve the "state-of-the-art" for open-source software with specific recommendations.
  - Published a report at the end of September: https://cpb-us-w2.wpmucdn.com/sites.gatech.edu/dist/a/2878/files/2022/10/OSSI-Final-Report.pdf
  - The most active discussions were in the memory-safe language area; how to drive adoption and strategies for supporting them
  - Incentive discussion led to an admission that there are a lot of things we don't know about open-source software and the practices.
    - We don't have a good understanding of what drives developers.
        - A lot of those folks are volunteering their time, and don't have much energy or resources to put into improving security
- **The Chair** commented that it seems they got the least results in the area that he thinks is the most critical, incentives for getting developers to adopt security practices. He asked for more information on what ideas could be tried and what the pushback was? Did they have access to the study done on developer work motivations that was done by the Harvard Business School and funded by the Linux Foundation a few years ago?
  - Mr. Keromytis replied that they are focused on the Linux Foundation and were involved in that report as part of the team.
  - An interesting finding is that money incentives for this group of developers likely won't work.
    - People that work on open-source projects are doing this because they believe in the goal of the project; they have some altruistic perspective.
    - So if we pull on that thread a little rather than giving them money, are there things we can provide that make their lives easier and remove the things they need to do but don't want to do?
  - Summary of findings:
    - There is a lack of understanding on how to build an effective human-in-the-loop ecosystem that creates secure software.
    - Given the multidisciplinary aspects of the problem, collaborative research among computer security researchers/experts and social, behavioral and economics (SBE) researchers is warranted.
    - Supportive tooling and metrics that provide transparency and situation awareness about the level of risk and exposure of OSS projects is needed and will likely support future research aiming at the socio-technical aspects and challenges in securing the OSS ecosystem.
- **Ms. Moussouris** asked if they thought about the potential of providing someone else to port critical open-source projects to memory safe languages? Original maintainers supervise the functionality
  - Mr. Keromytis replied that what she says makes sense if looking at this as a snapshot in time.
    - Maybe not rewrite the whole software, just the most critical parts.

        - He is in favor of this idea but there are issues:
            - Continuous maintenance. The person that makes this modification has to stick around.
                - As the project evolves, they have to take ownership of those parts
            - The person doing this would need a high level of specialized expertise
            - Lots of hidden costs here.
    - Ms. Moussouris agreed it would require having a specific group of maintainers using a shared library model of components written in memory-safe languages.
    - Mr. Keromytis added that a precondition is identifying the places where you get the most bang for the buck, identifying the most critical open source projects.
        - There is a list that the Open Source Software Foundation has been putting together that could be used as a starting point
        - Leads to the question of how to identify the projects that will become critical
        - The Chair wondered if Log4j would have been on that list and Mr. Keromytis replied that he didn't know.
            - In the software dependency management discussion, they concluded that they don't have the big map.
                - There are issues of origin: many of the projects find their way into products in a modified way so even an SBOM won't help.
                    - It's not maintained by the vendor
- The Chair asked about next steps
    - Mr. Keromytis replied one idea was to use GSA to start driving vendors to invest in the security of the ecosystems that have the open-source projects that they use; start putting in some resources and raising the bar
    - Mr. Scholl mentioned that they have been discussing the report with ONCD and OMB
        - OMB is drafting a series of actions and directions to consolidate the recommendations from the report into actions for agencies
        - Looking at integrating and coordinating with some of the other activities going on in government
        - NCCoE has been looking at ways to apply the SSDF to open-source software security

Software Supply Chain and DevOps Security Practices (Michael Ogata)

- Separate, tangential effort in open-source software at NCCoE
- Problem: Want to integrate DevSecOps into government operations
    - Many strategies already exist including the use of software security practices (SSP)
    - The Secure Software Development Framework (SSDF) has many references that can fulfill the practices and tasks that are identified with it.
- NCCoE will develop reference designs and implementations for software projects using commercially available products to demonstrate or improve capabilities.
- The SSDF:
    - A framework that divides the software development process into four large buckets:
        - Prepare the Organization
        - Protect the software
        - Produce well secured software
        - Respond to vulnerabilities
    - Practices are subdivided into various tasks and mapped to existing NIST references or business best practices from other secure development frameworks

- NCCoE Draft Project Plan
  - Went out in July for comment
  - Hoping to get final draft published beginning of November using all the responses
- Project Plan: Use Cases
  - Free and Open-Source Development
    - Will invite an open-source project into the center who wants to improve and demonstrate their security capabilities throughout their development process.
    - Instantiate that in NCCoE
  - Closed-source software development
    - Same process but with a closed-source, software as a service model
  - Will accept as many as they can accommodate for each of the use cases
- Project Plan: Overview
  - Invite developer organizations and various technology providers into NCCoE and build a system on prem with cloud aspects
  - Goal is to build a practice guide with step-by-step instructions
  - Hope to identify artifacts that can be shown to achieve the goals of the SSDF and identify gaps and improve automation between the steps
- NCCoE DevSecOps Workshop
  - 400+ attendees
  - Speakers from multiple companies
  - Recording slides will be available on https://www.nccoe.nist.gov/get-involved/attend-events/nccoe-devsecops-workshop
- Next Steps
  - Project description and Federal Register Notice to be published in November
  - Continue to build community of interest
  - Select participants for inclusion in the NCCoE use cases
  - Anticipate project kick-off in January 2023
- **The Chair** asked if they have gotten any interest in participating from open-source projects?
  - Mr. Ogata replied that they haven't started reaching out to projects yet.
  - The Chair requested to be kept informed.

National Academies Report on Cryptography and the Intelligence Community (IC): The Future of Encryption

Peter J. Weinberger, Google

Disclaimer

- Mr. Weinberger started by stating that in presenting this report, he doesn't represent Google, any government agency, the UN or anyone else.

Membership

- It's clear looking at the membership that this committee is about encryption and the intelligence community. Membership is made up of:
  - Intelligence community veterans
    - Two former general counsels of the NSA
  - Academic and professional cryptographers
  - System builders

- Computer security experts
- Hans Davies was their methodologist
  - Had an assigned methodology
  - No classified information in the report

Statement of Task

- Report is informed by the commercial world and the state of the commercial world
- Identify potential scenarios over the next 10 to 20 years for the balance between encryption and decryption and other data and communications protection and exploitation capabilities.
- Identify plausible scenarios and technology and other major drivers for explaining what's going on.
- Then assess the national security and intelligence implications and assess options for what the IC might do.
- Not asked to predict the future. They were asked to consider a variety of scenarios
- Did not ask for recommendations but they did describe common actions that the IC might take

Report Outline

- Report can be found at https://nap.nationalacademies.org/catalog/26168/cryptography-and-the-intelligence-community-the-future-of-encryption
- Report contains a lot of discussion not easily summarized

About Encryption

- He did not spend a lot of time explaining encryption given the expertise of the ISPAB but wanted to ensure a common vocabulary. For example,
  - Public key cryptography is used for things like key exchange.
    - A complicated cryptographic thing that goes on that assures your browser that you're actually talking to the National Academies website.
    - After you exchange information with the website, it downloads the report in encrypted form.
    - Two benefits:
      - No one could read the report (not a benefit in this case due to it being a public report)
      - No one can change the report on the way to your computer.
    - What happens after it gets to your computer has nothing to do with cryptography
- Applications of cryptography:
  - Confidentiality, integrity, and authentication
    - Verifying the authenticity of someone we're talking to
  - Emerging technologies:
    - Computing on encrypted data
      - MongoDB: you load your encrypted data into their database, you send an encrypted query, you get encrypted answers back,
      - presumably that database cannot read your data.
      - They can't see what you're doing
    - Passkeys – a way to avoid conventional passwords
    - zero-knowledge proofs…

Creating Scenarios

- The idea is to stretch your view of the future by looking at various scenarios of things that might happen as you look at a problem. They used the problem of quantum computers.
  - Look at all the factors that might contribute to whatever happens with cryptography

- Group them into categories and choose three categories (a.k.a. drivers)
- Look at the extremes on either end for each driver
- Choose some to study in detail

- Scenarios:
  - Research
  - Analysis
  - Synthesis
  - Storytelling & Role-play
  - Evaluation & Planning
- **Drivers of the Future of Encryption**
  - Three drivers that they chose to study are:
    - Scientific Advances
      - Two extreme points
        - Predictable
          - A predictable evolution of the science, whichever science is as being considered.
          - Quantum computer example for predictable:
            - Do not threaten current cryptography
            - Refined suite of quantum-resistant encryption algorithms
            - Specialized deployments of advanced techniques
        - Disruptive
          - The other extreme is surprises and disruptions of various sorts, and it'll be pretty much cryptography.
            - Quantum computer example for disruptive:
              - Unforeseen improvements in quantum computers
                - someone comes up with a quantum computer tomorrow
                - It's an era of misinformation, an adversary could attempt to persuade us that they have a quantum computer.
                  You don't need an actual quantum computer to get that disruption, all you must do is get important people in the US government to believe there's a strong possibility that an adversary has a quantum computer to get the disruption.
              - Significant cryptanalytic breakthroughs threaten deployed cryptography
              - Unexpected new applications of cryptography cause frenzy of attention
      - Cryptographic systems aren't used in isolation.
    - Society and Governance
      - Non-technical issue. Encompasses all the things that affect the environment in which the intelligence community has to deal with encryption and decryption
      - Includes changes in cultural attitudes toward privacy, trust and the extent to which the populace trusts institutions.
      - How nation states and individuals perceive risks and benefits of encryption and have designed standards, laws, and policies to address them.
      - Two extreme points
        - Fragmented example:
          - Governments demand local data storage and access, and limit privacy
          - Local technologies are the norm
          - Citizen mistrust poses internal and external challenges for IC

        - Can affect relationships with foreign governments
    - Global example:
        - Markets preserve a unified internet protected by strong encryption
        - International standards rule the internet
        - Citizen trust enables solutions for IC mission while preserving privacy
- Systems
    - Cryptography is embedded in bigger systems and the situation is that this situation in bigger systems is not altogether good.
    - Two extreme points
        - Chaotic example
            - Customers cope with the security they're delivered
            - Vendors deliver features and performance
            - Slow progress in security tools and processes
            - Security is a niche specialty for engineering and IT
        - The future that's least stressed and the future that we would like to see is what's called "mature." Mature example:
            - Customers demand secure products and services
            - Vendors view security as a must-have attribute
            - Market demand leads to significant progress
                - Burst of innovation
            - Security knowledge and expertise are both broad and deep
            - Protecting information is a lot easier
    - The way encryption technologies are implemented
    - Recommended reading – Ian Levy, chief technology officer at GCHQ, departing blog post: https://www.ncsc.gov.uk/blog-post/so-long-thanks-for-all-the-bits

- **Mr. Groman** asked if the current debate around backdoors and law enforcement mandating access and breaking into an encryption came into play at all?
    - Mr. Weinberger replied yes. That would fall under the society and governance driver. It's in the long collection of things governments might do to affect outcomes.
    - The Chair added that there's a multi-page text box on that subject in the report.
    - Mr. Groman asked if the report weighs in on the issue and Mr. Weinberger replied that the report isn't intended to.

Selected Scenarios

- There are 8 possible scenarios. They chose three to delve into:

Scenario		Scientific Advance		SocioGovernance		Systems
Scenario 1	=	PREDICTABLE	+	FRAGMENTED	+	MATURE
Scenario 2	=	DISRUPTIVE	+	FRAGMENTED	+	MATURE
Scenario 3	=	PREDICTABLE	+	GLOBAL	+	MATURE
Scenario 4	=	DISRUPTIVE	+	GLOBAL	+	MATURE
Scenario 5	=	PREDICTABLE	+	FRAGMENTED	+	CHAOTIC
Scenario 6	=	DISRUPTIVE	+	FRAGMENTED	+	CHAOTIC
Scenario 7	=	PREDICTABLE	+	GLOBAL	+	CHAOTIC
Scenario 8	=	DISRUPTIVE	+	GLOBAL	+	CHAOTIC

- For society and governance driver they always chose "fragmented" because it's more stressful
    - Believed it would give more useful information to the IC
- Variations in scientific advances and systems drivers exposed a range of future events

Scenario 2: A Brave and Expansive New World

- This is where breakthroughs in quantum computing is balanced with more secure systems and software and an orderly transition to post-quantum encryption
- International relations will dominate with a few major power blocks, each with its own encryption standards.
    - The committee felt that would imply that when you're communicating between blocks, you would use some kind of least common, and therefore hypothetically insecure, denominator.
    - This is a scenario in which the defense has the advantage, which means that the intelligence community would have to adapt to fewer technical means of gathering intelligence.

Scenario 5: The Known World, Only More So

- This one has predictable cryptography, the society and governance are still fragmented, and the systems are chaotic.
- Breaches remain common.
- Attacks on systems are the province of governments and private actors
- The offense has the advantage but we're all the victims of one offense or another.
- They're not predicting but this is the one that sounds like a prediction.
    - There have been big advances in system security but there have not been big advances in the outcomes of system security.
        - All those advances have just made the adversaries do more work.

Scenario 6: Colony Collapse

- This is a grim one but not inconceivable
- In this scenario, somebody breaks public key cryptography.
- Nobody quite knows what to do because there's not enough expertise around.
- The systems are bad so it's hard to patch.
- People don't trust their institutions.
- It's hard to patch due to fragmented standards. It's hard to patch and it's hard to do anything at all.

Key Finding Summaries

- There is more in the report. Here are a few of the findings:
    - Buggy systems are likely to undermine the security of systems that would otherwise be pretty well protected by good cryptography.
        - Not only software systems. There are also bad security practices in account recovery, bizarre password renewal procedures and requirements.
        - Every time you do anything that's hard for people that leads to mistakes
    - Fragmented societies are likely to make security hard for the organizations required to rely on encryption.
    - If you're going to make progress on encryption and security, you need technical talent which is in short supply
    - Computing on encrypted data has the potential to improve security and privacy for individuals and organizations

Implications for U.S. Intelligence

- Encryption is going to be important.
- It's an uncertain future, you need to plan for it.

- Implies you have to detect quickly which of these trends are going to happen.
  - It would be nice for the intelligence community to concurrently plan for alternate outcomes but that's extremely difficult.
    - There is no one intelligence community
    - What ICs do is driven by the national intelligence priorities framework
- There's a tendency in the government and the public for there to be panics about various things. Disinformation surprises need to be treated with some restraint and it's going to be up to the intelligence community to do that.
  - It is incumbent on the Intelligence Community to make clear to policy makers what is at stake
- Key Finding: With more adversary nations (especially China) seeking and making advances in encryption and as academic researchers (especially in Europe) continue to invest in cryptography, the advantage of the Intelligence Community will diminish if not disappear.
- May need to adapt personnel policies to accommodate short-term employees or external consultants to gain access to needed expertise

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

Update on Post Quantum Encryption and Cryptographic Transitions

Lily Chen, NIST

NIST Cryptographic Standards

- Public key-based standards
  - FIPS 186: Signature
  - 800-56 series: Key Establishment
- Symmetric key-based standards
  - FIPS 197: AES
  - 800-67: TDEA
  - 800-38 Series: Modes of Operations
  - FIPS 180: SHA-1/2
  - FIPS 202: SHA-3
  - FIPS 198: HMAC
  - 800-185: SHA3 derived functions (parallel hashing, KMAC, etc.)
- Guidelines (800-131A, 800-133, 800-57, 800-208)
- Tools (800-90 Series, 800-108, 800-135)

Quantum Impact

- Quantum computers have the potential to break current public-key cryptography such as RSA, Diffie-Hellman, and ECDSA.

NIST Post-Quantum Cryptography (PQC)

- Active research into PQC in last 10-15 years
  - Lattice-based
  - Code-based
  - Multivariate
  - Hash/Symmetric key-based signatures
  - Isogeny-based schemes

- Some of these date to the 70's

NIST PQC Scope

- Updating the public key cryptography standards

NIST PQC Process Update: Milestones and Timeline

- 2016
  - Crypto Conference
  - Announced call for proposals
- 2017
  - Received 82 submissions from 25 countries in 6 continents
  - 69 submissions satisfied the basic criteria as 1st round candidates
- 2018 - 2022
  - Held 2nd and 3rd rounds of submissions
  - 3rd round narrowed down to 7 finalists and 8 alternate candidates
  - July 2022 announced selection of 4 algorithms
  - 4th round candidates are key encapsulation mechanisms (KEM)
  - SIKE, has already been broken and needs to be replaced
    - New call for additional signature algorithms has gone out
- 2023
  - Planned release of draft standard for public comments
- 2024
  - Planned release of first version of standard

The 3rd round selection

- KEM – Crystals-Kyber
  - Module learning with errors (MLWE)-based
- Signatures
  - Crystals-Dilithium
    - MLWE-based
    - Fiat-Shamir signature
    - Relatively easy to implement
    - Uniform distribution
  - Falcon
    - Based on SIS over NTRU lattices
    - Hash and sign
    - Relatively complicated
  - SPHINCS+
    - Stateless hash-based signature – don't need to manage the state
    - Slower and larger signature than stateful
    - 800-208 specifies stateful hash-based signatures

The 4th round selection

- KEM
  - ClassicMcEliece
    - Proposed in 1978
    - Secure but may be tough for common usage

        - Very large public key
    - BIKE
        - Competitive performance among non-lattice based KEMs
    - HQC
        - Strong security assurances
    - SIKE
        - Eliminated because it has been broken
        - Based on isogenies of elliptic curves
        - Will continue to research isogeny-based cryptography
- NIST intends to select at least one additional KEM in the 4th round over the next 18-24 months

Call for Additional Signature Algorithms

- Interested in digital signature scheme not based on structured lattice
    - Want to diversify underlying security assumptions
    - Target specific applications such as those needing short signatures
    - Want mature designs
    - June 1, 2023, deadline

PQC Migration

- NCCoE initiated the Migration to PQC Project
    - Work with industry partners to look into real life applications of PQC algorithms
    - Currently have > 10 partners that participate in 2 weekly calls
- Working with standards organizations like IETF on migration issues
- Looking at hybrid solutions using current public key cryptography and PQC
    - This will be an application decision
    - NIST will try to accommodate if possible

Cryptographic Publication Review

- NIST has 45+ years of history publishing cryptographic standards
- 2021 NIST Cryptographic Technology Group established reviews for standards on a 5-year basis
    - Have reviewed 4 publications and 8 more are under review
    - Review board consists of 5 – 8 people
- Revising AES
    - Published 20 years
    - Adding clarifications and guidance on selection of keys

Cryptographic Transition

- Cryptography is constantly changing to become stronger
    - Moore's law – increased computing power
    - Quantum computers
    - Better cryptoanalysis techniques
- NIST Guided Transitions from
    - DES to triple DES to AES
    - SHA-1 to SHA-2 and SHA-3
    - RSA 1024 to 2048
- After 2023
    - Disallow 3key triple DES (800-67)

- Disallow PKCS1-v1_5 padding for RSA encryption (800-56B)

Mr. Venables asked about patent issues in PQC

- Mr. Scholl replied that there are no outstanding patent issues.

The Chair asked if there is a plan in place in case of a cryptanalytic break of RSA factorization

- Ms. Chen replied that is why PQC is urgent.
- The Chair reiterated that, with RSA, there doesn't seem to be an alternative and, if we did, cutting over to that alternative would take months or years. There's so much infrastructure that depends on RSA. It might make sense to have standards or certification on PQC with multiple algorithms where you have to demonstrate an ability to do both alternatives.
- Mr. Scholl agreed to continue discussions on this.

Final Board Reviews, Recommendations and Discussions

Steve Lipner, ISPAB Chair

Topics for Future Meetings

- Implementing the Risk Management Frameworks: What works and what doesn't
  - Suggested by the Chair
  - AI RMF, RMF, CSF, PF
  - Gaps between FISMA reporting, framework requirements, and agency actions
  - See discussion under "Board Actions"
- Continued discussion on the OMB Memo, M-22-18, Software Supply Chain Secure Development
  - Suggested by the Chair and Dr. Baker
  - Want more information on the potential for third-party certifications
  - (Maybe) How can they solve open-source software?
  - If possible, include an agency or vendor who can share their experience working with it.

Board Actions

- The Chair brought up the Risk Management Frameworks
  - AI RMF, RMF, CSF, PF
  - Ms. Miller indicated they may want to capture concerns about the magnitude of work put onto organizations with limited resources, who are already strapped in trying to define and manage risk. We need meaningful risk management discussions across the relevant communities (CIO, CISO, CPO) to determine the most efficient, effective way organizations can manage risk within current constraints. She indicated that she gave a note with some opening comments on this topic to Mr. Scholl.
  - Mr. Groman agreed and added that it would be good to eliminate inconsistencies.
  - Mr. Scholl acknowledged receiving the note from Ms. Miller and added that the discussion earlier included reaching beyond the CISO to the other C-level folks within the agency.
  - Mr. Groman agreed that this needs to be included in the mission planning. He also talked about the political problems such as difficulties with long-term planning coming from short-term appointees and agency secretaries taking ownership of the risk.
  - Ms. Moussouris raised the question of how to incentivize things that are seen as "thankless" tasks?

- Mr. Gattoni brought up budget issues; solutions can take multiple years, but budgets can be difficult to work with over multiple years. Figuring out how to give them some budgetary flexibility could help.
- The Chair mentioned that his impression is that complying with the CSF, PF, RMF and, eventually, the AI RMF where applicable, is hard to operationalize. He has heard from agencies that following up on these mandates from FISMA is a burden that is not always helpful in making their systems more secure. What can be done, in the context of these frameworks, that would be actionable and deliver results? He suggested that the 45-minute presentation that they had during this meeting wasn't enough and that the board needs more information on the dynamics across these frameworks and the process finishing with the compliance reports and/or the operational systems as configured.
- Mr. Scholl agreed.
- Ms. Flynn Goodwin asked if they had worked with the CIO Council and if these questions came up there.
- Mr. Scholl replied that they haven't worked with the CIO Council on these issues. They don't set the Council agendas, but they can suggest topics. He agreed they can do that engagement with the CIOs, CISOs and include the CIO Council.
- Ms. Flynn Goodwin also asked how they can reconcile gaps between what's reported by FISMA, what is required by the frameworks, what the federal government is doing, and how to close the gaps.
- Mr. Gattoni mentioned that CISA announced the release of the cyber performance goals. Part of that is intended to serve as a quick start guide for prioritization across the CSF for critical infrastructure owners and operators. He agreed with the Chair that it would be worth gathering more information, building up toward a letter that would get that effort in front of the board for the March agenda to see what they have come up with. That gives them time to socialize it and have a stakeholder engagement plan. We could get an idea of how CISA is promulgating this information.
- The Chair expressed a desire to hear from an agency that has either successfully or unsuccessfully implemented this and their lessons learned, all the way down to operational security. End-to-end, where it did or didn't work.
- Mr. Groman mentioned that it can be difficult to find folks able or willing to have that candid, open, and honest dialogue. Few CISOs and CIOs want to admit problems.
- **Suggested Action:** The Chair said he doesn't want to draft a letter at this point. He asked Dr. Baker, Ms. Miller, Mr. Scholl, and Mr. Gattoni if they were willing to suggest CIOs, CISO, or other individuals who are able and willing to talk about this, making sure they understand the ground rules and that this is a public meeting. We need folks who can be reasonably candid about what works and what doesn't work.

- **Letter in support of CSRB**
  - Mr. Gattoni suggested writing a letter in support of the work done by CSRB and its continued maturation. He liked the idea of modeling the CSRB work after the Crash Safety Board, where they don't wait for a crash before they start looking for vulnerabilities.
  - The Chair added that he would like to see them used for a real event, not just limit it to vulnerabilities.
  - Ms. Moussouris added that one of the biggest challenges for the CSRB is how to maintain a blame-free review standard and focus on what worked and what didn't in a non-judgmental way.

- **Action:** The Chair agreed to draft a letter and send it to Ms. Moussouris and Mr. Gattoni for review prior to sending it to the entire board for review.

Next Meeting

- The March 01-02, 2023 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 – October 26 and 27, 2022		
Last Name	First Name	Affiliation
Board Members in Attendance		
Lipner	Steve	SAFECode (Chairperson)
Baker	Brett	NARA
Fanti	Giulia	Carnegie Mellon University
Flynn Goodwin	Cristin	Microsoft
Gattoni	Brian	DHS
Groman	Marc	Privacy Consulting
Miller	Essye	Executive Business Management (EBM), LLC
Moussouris	Katie	Luta Security
Venables	Philip	Google
Board Members Not in Attendance		
Fitzgerald-McKay	Jessica	NSA
Hallawell	Arabella	WhiteSource
Maughan	Doug	NSF
NIST Staff		
Brewer	Jeff	NIST
Scholl	Matt	NIST
Carlson	Caron	HII
Salisbury	Warren	HII
McConnell	Andy	HII
Lurie	Kirk	HII
Speakers		
Romine	Chuck	ITL, NIST
Tabassi	Elham	NIST
Weitzel	Dave	MITRE
Snyder	Julie	MITRE
Sames	Christina	MITRE
Sonmez Turan	Meltem	NIST
Kelsey	John	NIST
Hall	Tim	NIST
Herckis	Mitch	OMB
Scholl	Matthew	NIST
Stine	Kevin	NIST
Temoshok	David	NIST
Galluzzo	Ryan	NIST
LaSalle	Connie	NIST
Deift	Abby	HQ DHS
Moussouris	Katie	Luta Security

Keromytis	Angelos	Georgia Institute of Technology
Ogata	Michael	NIST
Weinberger	Peter	Google
Chen	Lily	NIST
Registered Attendees		
Avery	Tory	Cybersecurity student
Boutin	Chad	NIST
Doubleday	Justin	Federal News Network
Doyle	Chris	Civil Air Patrol, U.S. Air Force Auxiliary
Doyle	Harry	HD Healthcare
Duplantis	Patricia	U.S. Government Publishing Office
Eggers	Matthew	U.S. Chamber of Commerce
Frascella	Chris	Epic
Friedman	Sara	IP News
Gardner	Zach	Keyhole Software
Geller	Eric	Politico
Guirreri	Joseph	USAF Retired
Heyman	Mat	
Irish	Shawn	Battelle
Kerman	Sara	NIST
Hykel	Daryl	Invitation Homes
Leithauser	Thomas	Wolters Kluwer
Lynch	Devin	Security Scorecard
Mace	Steve	NCTA
Nguyen	Thanh-Thien	Kaiser Permanente
Pascoe	Cheryl	NIST
Patrick	Jhamaal	Micolby
Pham	Peter	Mukilteo School District
Prado	Bernardo	Miami-Dave Aviation Department
Rogers	Susan	Sumitomo Mitsui Banking Corporation
Sakellariadis	John	Politico
Scarfone	Karen	NIST
Sedgewick	Adam	NIST
Sokol	Annie	NIST
Souppaya	Murugiah	NIST
Starkie	Edward	Kroll
Stoller	Travis	Wiley Law
Suarez	Jenn	Privacy Practice
Throneberry	Saundra	Lockheed Martin
Walther-Puri	Munish	Presearch Strategy
Weber	Bill	Cyber Foundry
Zigouris	James	FCC