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

July 13-14, 2022

Virtual Meeting Platform: Blue Jeans

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

Wednesday, July 13, 2022

Welcome and Opening Remarks

Steve Lipner, Chair, ISPAB, Executive Director, SAFECode

- The Chair opened the meeting at 10:08 a.m. ET and welcomed everyone to the meeting.
- Reviewed the agenda
- Reviewed mask/COVID policy

Board Member Introductions and Updates

Steve Lipner

- Executive Director of SAFECode
- Retired Director of software security at Microsoft.
- I've been working on a committee of the National Academies on the future of encryption for approximately two years. We're in our last review phase. I'm hoping we'll be able to brief that to the board at the October meeting.

Jessica Fitzgerald-McKay

- Chief of NSA Center for Cybersecurity Standards.
- News since last meeting: We have kicked off our Enduring Security Framework on threats to standards. We've had three meetings trying to get the community together to discuss current US posture and standards from a national security perspective.

Dr. Brett Baker

- Inspector General of National Archives.
- Previously at the Nuclear Regulatory Commission
- These are the things we're paying attention to in the IT community.
  - FISMA: We have an oversight role in that area. The schedule got tighter this year because the posting to OMB CyberScope is at the end of July. We have to move up our schedule for testing. I've been working with OMB, potentially going into August if we need to. They've granted most of the requests to do that.
    - ✓ The purpose of doing testing earlier and CyberScope reporting earlier is to use that for budget planning. The budget cycle in the federal government can be a two-year process, but September is very critical. Doing it before October 31 is a good thing but they didn't tell us until April.
  - Supply Chain: Supply chain testing is a big deal. We're trying to do more in that space.
  - Blockchain: MITRE's got some really important work that they're doing with Grant oversight, helping the efficiency of the payment request agencies, watching that reporting on those payment requests is okay. Grant recipients can benefit from using blockchain to reduce the reporting burden. From an IT perspective, we, including agency grant managers, can have a better sense of how money has been spent.
  - Quantum Computing: Looking at how quantum computing might have an impact on the blockchain technology and the encryption algorithms. It's immutable to a certain extent, but if those computers are cranking, we might have a problem.

Giulia Fanti

- From Carnegie Mellon University.
- Since the last meeting, I've been working on a report on the cybersecurity of centrally banked digital currencies, looking at how blockchains and distributed ledger technology can help or potentially hurt.

Brian Gattoni

- Chief Technology Officer at the Cybersecurity and Infrastructure Security Agency (CISA).
- Recently announced our partnership with NIST on our own post quantum initiatives to help critical infrastructure owners, and operators and owners of government networks, through their transition.
- I got back from Tel Aviv two weeks ago as part of the delegation for cyber week. We're doing exciting work with my counterparts in Israel's National Cyber Directorate (INCD) and in partnership with an organization known as the BIRD Foundation. Working on advancing investment in advanced analytics and automation of cyber analysis for operators both in CISA and INCD, even supporting the BIRD foundation setting up small businesses and startups in both the US and Israeli marketplaces.

Essye Miller

- Continuing my efforts on helping the next generation of cyber talent come to the forefront. Working with the National Cybersecurity Scholarship Foundation, to get high school kids to understand their depth and penchant for a career in cyber.

Arabella Hallawell

- Spent over 25 years in the cybersecurity space, working both on the technology research side as well as with various solution providers.
- Currently, I'm a chief marketing officer for an application security on demand company formerly known as WhiteSource that works well with organizations' security issues.
- Since the last meeting, I've spent a lot of time working with customers who are grappling with the current application security challenges as well as how to implement SBOMs.

Katie Moussouris

- Founder and CEO of Luta Security. Were a cybersecurity company that helps large organizations and governments deal with vulnerability disclosure and bug bounty programs.
- Since the last meeting, this is one of my three federal boards, and it is board season right now. One of the boards is a brand new one that CISA runs, the Cyber Safety Review board.
  - Created because of the executive order on cyber.
  - Completed our 90-day review of the log4J incident and delivered that report to the President
  - Hopefully there will be a public version of that report released. I'm looking forward to getting to that point.
  - I'm also involved in the ISTAC over in commerce and have been working on various export controls for technology.
    - ✓ It's heightened in the current conflict, the war, having to activate some other types of export controls.

Marc Groman

- Served in a second term with the Obama administration as President Obama's Senior Adviser on Privacy.
- Currently consulting and teaching cyber and data breach response at Georgetown Law School.
- Also served as an adviser to boards at the NSA and consult with the private sector and increasingly working with various stakeholders on Capitol Hill on privacy. The potential for monumental changes in privacy in the near future are huge. Hopefully this board and NIST are taking notice.

Welcome and ITL Update

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

My goal here is to take things up to the 60,000 foot level, talk a little bit about the context in which NIST operates, and give you some background on exciting things going on.

Cultivating Trust in IT and Metrology

- Partnerships
  - The entire purpose of the ITL is cultivating trust in information technology and metrology. And that word "trust" runs deep through everything that we do.
  - Trust is pivotal because we're often asked, "How is it you think you can make a difference with the small amount of funding that you have?" The answer is if we worked in isolation, we probably couldn't make much of a difference, but that's not the way we work. We work in partnership with industry. We work in partnership with other federal agencies. So $1 at NIST is multiplied a hundredfold, a thousandfold. That only works if we are a trusted entity.

New Positions

- New permanent Director, Dr. Laurie Locascio
  - Back after a 5-year hiatus
  - Confirmed by U.S. Senate on April 7, 2022
  - 17th NIST Director and second female NIST Director
  - Brings a deep understanding of NIST, as she was the Director of the Material Measurement Laboratory and, briefly, the Associate Director for Laboratory Programs.
- NIST Chief of Staff, Dr. Jason Boehm
  - Brings a deep understanding of NIST. Has served, over the last decade or so, as the Director of the Program Coordination Office

- ✓ Is a key position at NIST to ensure cross pollination across laboratories and other major programs at NIST. He was recently given job as the NIST Chief of Staff. Congratulations!

Quantum-Resistant Cryptographic Algorithms

- NIST announced 1st four quantum-resistant cryptographic algorithms
- This is the culmination of a process that started six years ago when Matt went to a cryptographic meeting in Europe to announce that we were going to embark on this journey, and he wanted to embark on this journey together with the entire community.
- Our cryptographers have worked diligently over the last six years with the cryptographic community, our partners, and stakeholders, and have come to the initial announcement of the first four cryptographic algorithms.
- Had an independent report assess the impact of our selection of AES more than 20 years ago.
  - Estimated economic impact was that this identification of AES added $250 billion to the economy over those two decades.

Cybersecurity Framework

- Started 1st major revision of the Cybersecurity Framework
- The Cybersecurity Framework started in 2012-13.
- There have been incremental revisions to the framework, but we got the clear signal from the community that a more major revision is warranted at this point.
- This is something that is of keen interest to the Deputy Secretary of Commerce, Deputy Secretary Graves.

EO 14028

- Executive Order 14028 on improving the nation's cybersecurity had numerous requests for NIST to engage in various activities.
- We released the summary report on our progress in implementing section 4.
- There was a summary report released two months ago on the labeling initiatives under the executive order as well as new guidance for supply chain risk management

National Software Reference Library (NSRL)

- Update to NSRL tools and library
- The NSRL is not a lending library. Through donations and purchases, we try to obtain the widest possible selection of released software as a catalog so that we can use hashing technologies to provide, particularly to law enforcement, the hashes of files that are expected to be benign.
  - This helps a law enforcement individual who seizes a computer to eliminate many, many files from any examination by doing the comparison of hashing.
  - Has the benefit of simplifying forensic analysis of machines that are seized through law enforcement activity or through national security activity.

Artificial Intelligence

- AI Risk Management Framework draft and workshop #2
- Most of the way through developing an AI Risk Management Framework.
- A robust risk management framework for artificial intelligence is not purely a technical issue. It is also a socio-technical issue because there are lots of sociological inputs or concerns that need to be addressed.
- A lot of the technical work that we're doing involves identifying and managing bias in AI.
  - Key driver for enhancing trust in artificial intelligence.
  - Very strong potential for consumers to have a loss of trust if we don't address bias in AI

  - Had symposium in April on identification and charting a path for responsible and inclusive AI. This was an important driver for the whole discussion.
- The National AI Advisory Committee, fondly called the NAIAC, held its inaugural meeting at the Department of Commerce on the fourth of May.
  - Meeting was led by five women, Laurie Locascio, the NIST Director and Under Secretary of Commerce for Standards and Technology, Gina Raimondo, Secretary of Commerce, Alondra Nelson, Acting Director of the White House Office of Science and Technology Policy, Lynn Parker, Director of the National AI Initiative Office, and Miriam Vogel, NAIAC Chair.
- AI and the Economy Symposium – April 27, 2022
- National AI Advisory Committee – May 4, 2022

Congressional Hearings or Briefings

- Securing DOTGOV: I was in a hearing on securing DOTGOV on May 17, 2022.
- Privacy in the Age of Biometrics: June 29, I had the opportunity to be a witness at a hearing in the House Committee on Science, Space, and Technology Subcommittee on investigation and oversight of privacy in the age of biometrics.
- On July 13, there was a briefing on CBP's use of biometrics and facial recognition for the members of the Homeland Security Subcommittee on Border Security, Facilitation, and Operation.

Budget Initiatives

- We did get a certain amount of additional funding in FY 22.
  - About $5.25M for cybersecurity, including:
    - ✓ A small bump up of almost 600k for artificial intelligence,
    - ✓ An additional couple of million for Quantum Information Science,
    - ✓ A small amount for standards for critical and emerging technologies, and
    - ✓ A tiny increment for NIST's work in diversity, equity and inclusion.
- Our requests for FY 23
  - We are requesting an additional:
    - ✓ $18 million in cybersecurity,
    - ✓ $15M for Artificial Intelligence and for Quantum Information Science,
    - ✓ $8M for Standards and Critical Emerging Technologies, and
    - ✓ $6M for NIST Workforce Diversity,
    - ✓ The request has been supported by the Commerce Department and looking forward to resolution of any questions that Congress may have.
- The money for cybersecurity includes money for privacy.

Celebrations

- 50 years of Cybersecurity research
- This year, we celebrated the 75th anniversary of the establishment of applied math and statistics as an entity at NIST
  - The entity for applied math and statistics, which has changed its name over the years and has two divisions in my laboratory, started 75 years ago
  - So we had a symposium to celebrate that.
- We are celebrating 50 years of cybersecurity research at NIST,
- Next year we celebrate 60 years of biometrics research at NIST
  - We started in partnership with the FBI on electronic representation of fingerprints back 60 years ago

Discussion

- **Ms. Moussouris:** Two questions, one of them about the hash library. What hashing algorithm?
- **Mr. Scholl:** They use multiple hashes, even MD5. Most of SHA-1, SHA-2, different variants of SHA-3, different variants of SHA-2 as well. What the forensic scientists in the NSRL have found is many of the local, state, and county law enforcement who might want to use NSRL are sometimes still stuck on some legacy implementations. They have asked us to continue to support that even though there is a potential of collisions with MD5. The legal issues that might or might not be created with the use of an MD5 hash in a forensic case is not our purview.
- **Ms. Moussouris:** My second question: on your budget slide, was the amount for DEI correct with the decimal point where it was?
- **Mr. Romine:** That isn't the total amount. That's the additional amount that we're providing.
- **Ms. Fitzgerald-McKay:** Back to the NSRL, is that a new capability or a highlight of a feature that's always existed?
- **Mr. Romine:** We've had NSRL in existence for close to a couple of decades. This is a new expansion and some new capabilities that we are providing.

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

NIST Cybersecurity and Privacy Update

Matthew Scholl, NIST
Kevin Stine, NIST

Clarification on Budget

Mr. Romine talked a little bit about FY 22, and FY 23 budget cycle. To provide a little bit more context on those numbers:

- We anticipate the additional $5.25M FY 22 funding showing up for our use soon.
- Coming in late in the fiscal year,
- Will be applying those resources to support some critical need areas and allow or initiate modest expansions in some priority areas, including our cryptography program, privacy, identity and access management, and supporting improvements to the infrastructure in the application for the National Vulnerability Database.
- Two specific directions:
  - One was an increase of $750k to apply to NCCoE genomic cybersecurity work. This is in addition to the last two fiscal years' $1.25M for that same program.
  - The other specific direction was $500k to support the initiation, under the NICE programs, of the Regional Alliances and Multi-stakeholder Partnerships Stimulation (RAMPS) program to improve cybersecurity workforce.
- We ran a pilot of the RAMPS program in 2016.
  - We issued five grants to different regional alliances around the country in the amount of about 200k each to help pull together local or regional resources, local government, local nonprofits, local industry, and community to really understand the local or regional cybersecurity workforce needs of those communities.
  - The funding that we provide and, in some cases, matching funding from others in the communities are used to help spur improvements to the local cybersecurity workforce.
  - That pilot model received a lot of positive attention, including through the Security Solarium Commission, and subsequently through NDAA
  - Part of our FY 23 request includes additional resources beyond the $500k to support continued growth in that area.

- The FY 23 request from the administration reflects an $18 M increase for our cybersecurity and privacy program.
- We intend to continue expansion that will allow us to do more high impact work in key areas including work in privacy, identity, supply chain, and workforce in addition to some other key areas called out specifically in the request narrative that's come out from the administration.
- Always open to getting the Board's input and ideas on budget for future years and on those key areas.

The Chair had a question on funding: Since you mentioned getting that incremental funding late in the year, how much does that limit you? What are you able to do or what are you not able to do that you'd want to do with it?

Mr. Stine replied: It does limit options because that funding must be used within the fiscal year that ends September 30. Sometimes, we initiate acquisition processes in key areas before we get funding. When funding shows up, we're able to execute those much more quickly. That is an option we're using here. In other areas, we can put funding on larger contract vehicles or grant opportunities that may be already out there to leverage that funding in a useful way for the program. That might present some implementation challenges regarding how to use the funding now but that becomes the starting point for our FY 23 budget for next year. Even if you're under a continuing resolution, that becomes the new bottom line. That allows us to make more informed planning decisions so when we roll into FY 23, on October 1, we have a better idea of where our resourcing will be and how to make planned improvements there.

NIST Post-Quantum Cryptography (PQC)

- On July 5, 2022, we announced the initial selection of the first sets of post quantum algorithms that NIST was standardizing.
- Tomorrow afternoon, I (Matt) will be up here with Natasha Cohen from CISA. Natasha and I, are two reps of agencies that are working on National Security Memo 10 (NSM 10), prepping the US government and the Nation for a quantum transition.
- We'll go through the work that we're doing jointly in addressing NSM 10.
  - This is the first set of quantum algorithms but certainly not our last set.
  - We still must standardize them. There is still a significant and important block of work ahead of us.
  - We have four algorithms identified.
  - Two primary algorithms:
    - ✓ a primary key establishment mechanism that's going to replace our Diffie Hellman algorithms (CRYSTALS-KYBER), and
    - ✓ a primary signature algorithm (CRYSTALS-Dilithium) that will help us to replace RSA.
  - We will then generate the first set of standards, working with the submission team and the public community to define the specifics of those implementations, so that they meet the requirements that NIST laid out in our initial Federal Register back in 2016.
- Security strengths of those implementations.
  - We will decide on the implementation parameters:
    - ✓ Which will allow to be built in commercial products with enough specificity that it can be interoperable.
    - ✓ To ensure correctness in its implementation.
  - We've identified:
    - ✓ An initial key establishment method,
    - ✓ An initial signature, and
    - ✓ Three other signatures.
      i. One for a place where there might be some uniqueness needed for key sizes.

ii. Two others that potentially will work with other kinds of edge use cases or are not necessarily based on the mathematical structures of the primary one/
    - We wanted to have a little bit of genetic diversity in the mathematics of the signatures, so, should there be a quantum breakthrough against a lattice structure, we have an initial immediate backup.
- After we get all these standards done, one of the deeper thoughts is to look at what other mathematical diversities we should bring into our portfolio so that we have a little more resilience against what might emerge in the quantum information sciences against those types of mathematics.
- We anticipate this process of finalizing these standards around 2024. The finalization of the standards will trigger a lot of different events.
- Names: The current names, CRYSTALS-KYBER, CRYSTALS-Dilithium, FALCON, SPHINCS+ are the names that the submission teams provided to us. Those names are likely to change.

Cryptographic turbulence

- Many other cryptographic transitions are going to occur between today and the time we transition into PQC.
- Next 5 to 10 years will be a time of cryptographic turbulence. In addition to the PQC standards, we are going to deprecate and disallow any use of three key triple DES starting this December. We are updating our signature standard.
- Will have some cryptographic changes to our current suite of signature standards
    - We are going to be pulling things from inventories.
    - We're also going to be looking at what industry is using commercially, has been using historically safely and securely and whether we should also be including that even though they may or may not be quite as safe going forward as well.
    - Working very closely with the privacy engineering team and looking at the use of encryption technologies in privacy enhancing cryptography, looking at using some of our current and quantum safe algorithms in new modes that are specific to privacy enhancing capabilities.

Light weight block ciphers that are still as strong as AES 128

- We continue to work on lightweight block ciphers. Authenticated encryption with associated data, for use in some of the new, smaller, resource constrained, small CPU size, small memory size, small bandwidth size application that still need to be as strong as an AES 128 implementation.
- So even though it's lightweight, the lightweight is going to apply to computation, bandwidth, or circuit size, but not the strength of the block cipher. That will also be coming soon in standards.

Need Feedback

- We have been asking a lot of the encryption community to review our work, to provide us with inputs, to check, and give us feedback on our implementation parameters.
- Will continue to need that going forward.

Mr. Romine asked for an estimate of the number of cryptographers NIST had 10 or 15 years ago, compared to today.

Mr. Scholl replied: NIST has 31 cryptographers as well as a set of guest researchers who work with us from around the world including France, Belgium, Turkey, South Korea. and Japan. It was probably a little less than half that 10 or 15 years ago. In the last 10 years, we've almost doubled the staff working on encryption and cryptographic technologies. That's mostly due to the work of Dr. Lilly Chen, who runs the Cryptographic Technology Group and her phenomenal outreach, both to folks in the professional

associations and in academia. We have been able to be competitive in hiring those folks and keeping them, considering it's such a desired field.

Ms. Moussouris asked for the commercial implementation guidance:

- What work is NIST doing around the sources of entropy? That was one big area where commercial implementations failed to implement the standards even though the standards were very clear.
- What work is NIST doing around cryptographic algorithm agility in commercial implementations? That is another one where we've seen cascading failures in commercial implementations and the inability to effectively patch a lot of these issues.

Ms. Moussouris clarified that cryptographic agility is the ability for you to swap out an algorithm that is no longer safe. This is something where we've seen hard coded encryption algorithms in software, making it incredibly difficult to upgrade when a problem is found with the core algorithm that's used.

Mr. Scholl replied to the first two questions.

- **Agility -** Crypto agility is a significant issue in commercial products and in the US government IT infrastructures that have long legacy implementations. Many of them are burned in silicon, especially when it comes to encryption technologies.
- In the commercial space, we had a long history of incentivizing implementation of encryption in hardware for speed of performance, especially implementation of primitives in hardware, which is the opposite of agility.
- Agility was one of the PQC selection criteria we looked at, which had to do with size compatibility of:
  - Gate equivalency, like how much memory an implementation might take,
  - key size compatibility,
  - and memory compatibility. for example, how easy is this to implement in a TLS implementation or is it not, on the assumption that the TLS implementation is in software to begin with?
- One of the concerns that the White House has about doing this transition is understanding the potential lack of agility.
- Inside the government, we are identifying where our legacy systems are, where hard coded hardwired systems that are burned into silicon are, and where the difficult places are, that don't necessarily facilitate an update such as space systems, where you can't just have someone walk in and push an update.
- We are relying a lot on the industry side to help.
  - A challenging opportunity. Example: border proxies that need line speed and security are hardwired in silicon rather than software and they don't update. There might be an opportunity for industry to resell to the government, which they might like.
  - One of the things that's going to have to happen with the NSM 10 work is for government to identify where those expenses are, where we don't have the agility, where we're going to have to do some of these repurchases and start to plan and budget for that as well.
- **Entropy -** we recently started to require entropy validations for encryption that is going to be used by the US government. Prior to that, entropy requirements were just a description. They told us what their entropy sources were, if and how they were conditioning their entropy, and where it was coming from.
  - Have Black-box bench entropy testing that can tell us your entropy is bad. It's difficult in any level of assurance to say your entropy is good, but I can tell you if it's bad.
  - Provided Black-box testing to our conformance laboratories, so that we can audit/run a test to ensure that the implementations are adhering to 800-90B, our entropy requirements.
- On the research side:

- Every minute we publish out a new random source.
- Working mostly with our physics laboratory to see if there are actual, true quantum random sources that we can incorporate into that.
- Mixing in photonic sources of entropy. These are like light bouncing on a piece of tinfoil for noise. So even though they call it quantum, it's not what we would call quantum.

The Chair asked: How prevalent is FIPS 140 certification in the commercial industry and how prevalent is actual use of FIPS encryption, and the use of the FIPS modes?

Mr. Scholl replied: The prevalence of FIPS 140, our conformance testing requirements, is ubiquitous in the government. It enforces the use of NIST approved encryption algorithms as the primary implementation in a commercial stack. Being used in FIPS mode is a different question. Often, when a product gets into an infrastructure and it is configured, it is not uncommon for it to be configured and used not in FIPS mode. The usability of that implementation is not as good as it should be.

90B testing:

- We started 90 B testing in November.
- Also transitioned from FIPS 140-2 to FIPS 140-3, which has some different sets of requirements.
- Transitioned some key lengths and some key strengths and have the triple DES transition coming up.
- Resulted in a large surge of participants in the FIPS 140 program:
  - Submitting implementations before the deadlines hit or
  - Submitting implementations after because they want to be some of the first ones in.
  - The growth of submissions was significantly more than 100% from the previous years.
  - Currently behind in clearing those out from our side of the validation. It's a focus area to ensure we test them so that they're good, and we're able to get them into our infrastructure. Kevin, Chuck, and Matt have been working jointly with the CMVP team and the Canadians to look at ways to streamline and where, like in entropy, we should not be taking risk, versus where there is potentially lower risk, requiring less oversight.
  - This is a possible discussion are with the board going forward.

The Chair added that FIPS Mode is another possible future topic.

Ms. Moussouris: There's a lot of attention on going dark and the ability for law enforcement and others to identify Child Sexual Abuse material and other material that otherwise would be shrouded in encryption. Is there any hope with these new encryption algorithms to provide, potentially, a way to search for these without decrypting?

Mr. Scholl: We have a lot of discussions about this issue and the difficulties it presents. We work on the technical standards. Our mission is to make those technical standards as strong as we can practically make them. They are not designed or built with lawful access as one of the selection criteria. It is an interesting concept to look at future aspects of potential partial or even fully homomorphic encryption technologies and if that might be a possibility in future implementations. Regarding our standards and our cryptographic algorithms and primitives that we create, we make them as strong as we practically can. The access is a policy issue that needs to be addressed by agreements between the government and the commercial products that are implemented.

Ms. Moussouris clarified her question: I am very much against having any backdoors built into the encryption algorithms. Is there any work being done to explore the feasibility and entertain proposals for the purpose of identifying the bad material without requiring decryption or a backdoor?

Mr. Scholl stated: We are looking at, and continue to conduct research in, homomorphic encryption, which allows for search into encrypted materials without decrypting the material regardless of what the material is. For us right now, that is still a research program. We're still determining if there are existing primitives in our inventories that we could use for it without having to create new ones and there's still a lot of debate around the practicality of it.

Mr. Stine replied: Tying this back to our budget discussion, part of the resourcing for both FY22 and in the requests in FY23 is to increase our capabilities to conduct more this type of research; that convergence of our privacy needs, privacy enhancing technologies, privacy enhancing cryptographic capabilities, to understand more about those challenge spaces and contribute to the research domains there.

On agility: The agility discussion is a segue into some of the work we're doing at the NCCoE around helping organizations prepare for their cryptographic transitions and there's a FIPS 140 threat that we can weave in there as well.

- We have launched a "how to prepare for post quantum migrations" project.
  - Collaborators include CISA and others across the industry and the USG.
  - How do we help organizations prepare to put those algorithms in action?
    - ✓ Doing the legwork now to understand the criticality of their data. Where they're using encryption today, what are the protection needs of that data?
    - ✓ And then helping them to determine and manage those risks and to prioritize their use of cryptography based on the protection needs of that data.
  - Intend to produce blueprints and guidance to help organizations, whether it's critical infrastructure owners and operators or state and local or federal agencies, or other users, to begin to prepare for that transition.
  - Agility is a key point as this is not the last transition we will have. Can we put in place the right tools and resources and blueprints to help organizations establish these processes, for this and future transitions?

FIPS 140 CMVP Validation Process

- As these algorithms are used and productized, there will be a likely path into our FIPS 140, CMVP validation process.
- There's a challenge there and an opportunity for us to take a hard look at how that program is executed. How we do engage with the labs, do the testing, and engage with the vendors?
- We have launched a project at the Center about automating different parts of the CMVP program and maintain that high level of assurance in more efficient and more effective ways.

Cybersecurity Framework

- Cybersecurity Framework Update – this will be presented later in detail

EO 14028

- We submitted two reports to the White House over the last couple of months:
  - Report on Section 4(w):
    - ✓ Summary of the cybersecurity labeling pilots for both consumer IoT and consumer software,
    - ✓ Identified opportunities and considerations for next steps in those labeling programs.
  - Report on Section 4(x):

  - ✓ A summary of the progress made for all the different activities that were directed in section 4 of the EO, enhancing software supply chain security,
  - ✓ That includes some of our activities that have occurred since our last meeting around, for example, the final publication of our cybersecurity and supply chains guidance 800-161.
- We've issued an update to the secure software development framework (SSDF).
  - There was a good convergence of our own interest with the executive order.
  - The SSDF update gives us a great opportunity to pivot that into producing tools and resources to help organizations most effectively apply the SSDF, use the practices in practical ways.
  - Will be announcing the draft project description at the NCCoE, as well as a workshop to begin to look at the SSDF practices, help identify and prioritize some of those, both in the context of commercial as well as open-source software, and work with the broader community to demonstrate the application of some of those practices in different scenarios and development environments.

Position Navigation and Timing (PNT)

- There was an executive order in the previous administration on position navigation and timing (PNT) that directed NIST to develop a cybersecurity framework profile for the use of PNT.
- We did that and we issued that approximately two years ago in response to that executive order.
- Recently issued a draft update to that profile that includes increased incorporation of cybersecurity framework categories and subcategories where, through application of the profile, folks have identified opportunities to put even more of the framework practices, outcomes and informative references into that profile in the context of responsible use of PNT.
- Comment period that closes August 12.

Space-based Operations

- Asked to provide some assistance in areas of space-based operations and commercial space-based operations.
- This is a new vertical for us.
- We worked with the National Space Council through space policy directive 5, and the previous administration's cyber strategy that asked NIST to provide some guidance to commercial space operators for cybersecurity.
- We published a report on considerations for commercial space operations.
  - We had the PNT document, which was focused on the services coming from a space-based operation, and another NIST IR that focused on the space vehicle.
  - Worked with the DoD on writing and publishing a report and some guidance on ground station operations and security of ground station operations.
- Our work has been in support of and in coordination with offices like NOAA, the Commerce's Office of Space Commerce, the DoD, with NASA, FAA, FCC, at the GPS executive committee, and the people who are space mission leads. Currently working with them in providing cybersecurity support as part of the mission areas.

SP 800-53

- We published a new online version of 800-53 in a data wiki format.
  - Allows users to give immediate feedback to a specific control, suggest a modification, suggest a new control, add on to someone's suggestion, and provide a more dynamic and active ongoing feedback rather than waiting for us to put out an open call for comment every two or three years.

- We will keep this running and public. You will be able to see comments as they move from submitted, to adjudicated, to going to be implemented at this date, and then pulled into the current version.
- These documents are also going to be able to be reflected out through a new data format, the cybersecurity and privacy reference toolset that Kevin, Murugiah, and other folks within NIST and within my teams have been putting together, so that we publish our documents as content rather than as PDFs.
- Control catalog will be available in multiple formats: PDF, XML, comma delimited, JSON, or whatever is the easiest and most usable for your tools to adjust and use as well.
- Helps the user to be able to sort, search, and then modify and customize for themselves as well.

SP 800-171

- At end of July, we will also be opening the 800-171 series for a pre-call for comment as we look at putting in revisions for that.

Cybersecurity Education and Workforce

- This is a priority for the Office of the National Cyber director, Chris Inglis.
- NICE:
  - We had our first in person NICE conference in early June, with approximately 500 attendees. That conference was followed by a symposium for the National Centers of Academic Excellence for Cybersecurity, managed by NSA and CISA.
  - Next year's NICE conference is going to be hosted in Seattle. Expected dates will be in early June.
- US Cyber Team
  - First ever US Cyber Team, think US Olympic team for cybersecurity, participating in international competitions. Funded by NIST in part along with others in the industry as well as in government, including CISA competed in the International Cybersecurity Challenge in Greece in early June and finished third.
- Last week, the second year of the US Cyber Games was launched beginning with a Capture the Flag event, a US Cyber combine and will culminate in a drafting of the next US cyber team on October 17, 2022. The U.S. will host the next international competition next summer.

Cybersecurity Apprenticeship Sprint

- Next week, Department of Commerce with NICE, as well as the Department of Labor, are going to announce a cybersecurity apprenticeship sprint. The purpose of this sprint is it's going to be an opportunity to raise awareness and really accelerated use of apprenticeships as a pathway for employers to address their cybersecurity workforce needs.

NCCoE Director Position

- Jeff Greene – NCCoE director, and former ISPAB member left NIST at the end of June. Mr. Greene was detailed to the White House National Security Council for 18 months. He will still be working in cybersecurity space and will continue to work with NIST. NIST will open a formal vacancy announcement for the next NCCoE director soon.
- Natalia Martin – Acting Director of the NCCoE
  - While Jeff was on detail to the White House, Natalia Martin has done an outstanding job serving as the Acting Director of the NCCoE, in addition to her deputy NCCoE director role.

AI discussion – societal impact

- **Mr Stine stated and asked the Board:** We continue to receive interest and requests, not just on our technical guidance, but on how do we incorporate and consider more societal impact in our technical guidance.
  - Certainly that's the case of AI. That's absolutely the case in privacy. We're seeing it more in our genomics work. And I think over the last several months it's been a very popular topic in the context of digital identity. Particularly in the case of where identity and facial recognition intersect and even, to a lesser extent, other biometric types.
  - We produce very strong technical guidance, but we recognize that that technical guidance must be viewed in the context of, not just the organizational need, but also societal impact.
  - Would like the Board's perspective on opportunities you see, challenges you see, others we should talk with about how to reflect those important societal needs and attributes most effectively in the technical guidance.

Ms. Fitzgerald-McKay asked: Is the idea that you would have societal considerations embedded in your publications in the way standards have security considerations or is this a new thread of work entirely for us to broaden the scope and impact? or a little bit of both?

Mr. Stine replied: It certainly could be both. Right now I'm thinking in the context of how do we incorporate more societal need considerations into our technical guidance? Where does it make sense? Where is it appropriate? How do we most effectively do that? But this certainly could be a much broader area that we want to look at more comprehensively across the program.

The Chair commented: I'd suggest that we queue up a preliminary discussion on the societal impact topic for the board discussion, either this afternoon or tomorrow afternoon. That's certainly something we ought to look at.

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

NIST Updates to the Cybersecurity Framework

Cherilyn Pascoe, NIST
Adam Sedgewick, NIST

Introductions

Cheri Pascoe

- Joined NIST September of last year
- Second time presenting before ISPAB. First time was six or seven years ago, in my prior job as a congressional employee working for the Senate Commerce Committee.

Adam Sedgewick

- I guess relative to NIST, I'm a new employee. I've been there for about 10 years.

Cybersecurity Framework 2.0

- Updates have begun. To participate, go to: https://www.nist.gov/cyberframework
- Background:
  - The framework is widely used by both private and public sectors within the United States as well as around the world to help organizations manage and reduce their cybersecurity risks.
  - The cybersecurity framework was always intended to be a living document. It was intended to be refined and improved and evolve over time as risks and technology changes.

- Look forward to hearing your thoughts on the framework and things that we should be considering as we proceed with this update.
- Recent RFI in February provided feedback from stakeholders and various themes that will be driving our work associated with this update.
- We are looking at making bigger changes in the framework than last time. This will be CSF 2.0, not 1.2.
- The framework is composed of three different pieces, the core, the profile, and the tiers. We are looking at changes across all components of the framework.

- History and timeline of the framework.
  - First developed in 2014, directed by executive order.
  - Was later codified by Congress.
  - Updated in 2018 with version 1.1.
  - We don't have an executive order or Congress driving our timeline and are able to work with stakeholders and have them drive the extent of the changes and the timeline associated with this update.
- RFI
  - Issued in February
  - This sought feedback from stakeholders about all our cybersecurity resources and the direction that we're going with a program.
  - Specifically asked about the Cybersecurity Framework, how folks are using it, changes that they would like to see, what's working, what's not, and how organizations are using the framework in coordination with other resources.
  - NIST has a lot of cybersecurity resources, a number of new frameworks like the privacy framework, secure software development framework, our workforce framework, our new IoT cybersecurity work, our supply chain cybersecurity work, all of which has largely been developed since the last update of the framework.
    - ✓ We want to know how organizations are using these resources together and how we can improve the alignment and interoperability between these resources and make it a little easier for folks.
  - Also asked specifically about supply chain cybersecurity.
    - ✓ Anticipate this to be a big component for the update to the framework.
    - ✓ RFI also asked more broadly about how we can address this issue at NIST.
- RFI Analysis
  - In July, we released a high level summary of the RFI responses.
  - Received about 135 comments from organizations, including dozens of associations that represent thousands of organizations.
  - Happy with the feedback that we received and the high level of stakeholder engagement.
    - ✓ Really high levels from financial and healthcare sectors
    - ✓ Some level of engagement from government and academia. We will seek to engage those stakeholders directly and see if we can fill in the gaps there.
  - Responses fell into 7 themes.
    - ✓ We are organizing the updates to the framework by themes identified in the RFI analysis and will be using these themes in discussions with stakeholders in breakout sessions.
  - Theme 1 – Focus on maintaining and building on the key attributes of the CSF with the update.
    - ✓ Had many glowing comments on the framework and its usefulness.
    - ✓ Many comments on keeping the CSF key attributes in the update.
    - ✓ Key drivers:
      i. We will be cognizant of the need to avoid changes that would limit its widespread use,
      ii. Focus on maintaining flexibility,
      iii. Keeping it usable for different audiences.

- Theme 2 – Align the CSF with existing efforts by NIST and others
  - ✓ We specifically asked about this, but commentators did agree that we do want you to align the cybersecurity framework with existing NIST resources, whether that be the RMF, the Privacy Framework, the Secure Software Development Framework, or Workforce Framework,
  - ✓ Also align with non NIST resources:
    i. Resources from other federal agencies that have regulatory considerations in cybersecurity to increase alignment of regulations associated with the framework,
    ii. International alignment was also a big topic. We want to make sure that the framework is not just a US product. It's something that's used around the world. We've been doing a lot of internationally outreach.
- Theme 3: Offer more guidance for implementing the CSF.

- Had significant requests for more guidance, including the desire for greater detail in the Cybersecurity Framework, of course while maintaining a non-prescriptive technical approach. A big challenge is how do we keep it short? How do we keep it simple how do we keep it flexible, while answering the call from stakeholders about the need for more guidance to implement the framework?

Mr. Groman commented: Clear and simple are very different. We need it to be clear and understandable, but it won't be simple. It cannot alleviate the need for somebody on the other end to think and it never will. The idea of expanded guidance is great, but the clear, simple align with everything is going to be a challenge. I'm working on a privacy law that's been moving right now in Congress and the definitions are new again. But that said, it's an extraordinary document.
Externalities of Cybersecurity This may not be a CSF question. This may be more for the RMF. One of the things from my perspective that is not baked into the RMF, or this is, what I would call for companies, the externalities of cybersecurity. In other words, when we talk about risk here, and you're a company and you're assessing your cybersecurity posture, you'll identify risks to the company and then we compare it to how much it's going to cost me and reduce my bottom line. I don't see where in the RMF, we are asking companies in the United States to contemplate the risk of their bad security to national security or to issues larger than your company's fourth quarter earnings. I feel that, as we move forward as a nation, as the threats are evolving, as our cyber challenges are evolving, those things have to be here somewhere. Part of cybersecurity is risk assessment, identifying your cybersecurity risks, taking steps to mitigate risk, and accepting your residual risk. We need to tell companies or entities or public entities, as you assess risk, like look at these other risks to national security that you're not contemplating. So does that fit into this or not?

Ms. Pascoe replied: I think that's a conversation we want to have with stakeholders. To your question, we're at the beginning of the process so we don't have the answers, but I think what you're raising is important.

Mr. Sedgewick added: A lot of your questions pertain to the definition of critical infrastructure and what it means to do cybersecurity if you're a critical infrastructure company:

- There are some elements of the Cybersecurity Framework specific to critical infrastructure.
- Some of the recommendations relate to how you receive information and how you share it.
- You need to be thinking about how you engage with government.
- Stakeholders have indicated a desire to retain "critical infrastructure" in the CSF.

Ms. Moussouris asked two questions:

- **Maturity Models:** Theme 1.2 was the voluntary flexible model being appreciated. Does this also fit into theme three on enhanced guidance for implementation? Is any work being done on creating a maturity model to ease the adoption of the NIST CSF and potentially to deal with the backwards compatibility and adoption problems?

- **ISO Publication:** Have you looked at submitting this or maybe 2.0 to ISO to be codified as an international standard? As an ISO editor, I know there's a couple of tricks to getting an ISO standard to be offered free of charge forever, and that's previous publication. That would drive that goal of having it be a truly adopted international standard.

Ms. Pascoe replied: Measurement and maturity. models came up a lot and are part of theme 5. Something we're going to be talking about is whether that's appropriate. Small and medium sized businesses need more guidance, but are we at a point where we can develop a baseline or are there subcategories within the cybersecurity framework that are more important than others? That's something that we're going to talk about. I think we want to maintain the risk-based approach that's in the framework and teach organizations how to bake the cake and not give them the cake.

Mr. Sedgewick added: That the tiers are cousins of maturity models and there has been some work on looking at maturity. Models, like the energy sector maturity model, and showing how that can align with the CSF. There is some foundational work. We could do some work thinking about how to align them better.

There is work going on in ISO currently. It was industry that took it to ISO. Much of the initial planning was about transitioning the CSF entirely into an international standards body. Stakeholders overwhelmingly said don't do that. Everyone, except one, felt much more comfortable with us maintaining it. I think we have a good model now where we maintain it and we have work going on at ISO at the same time. We're happy to connect with you to talk more about what you're seeing in ISO and those types of strategies.

Ms. Miller asked a follow-on question: (re: Mr. Groman's comments about externalization of risks, national security, and the government). We must be conscious of the information sharing requirements we're expecting of our industry partners. How do information sharing or information sharing requirements play into the update?

Mr. Sedgewick replied: I don't know if that came up during comments. I do think we will consider how an organization processes, receives, and ensures information. Making sure we can do that as easily as possible or help organizations.

Ms. Pascoe broadened the question: We heard a lot from organizations that are subject to existing regulatory requirements, or potential future regulatory requirements like the SEC reporting and the CISA incident reporting rules. That is something we have to consider within that broader alignment bucket. We also must decide what we mean by alignment. As organizations develop a profile, they need to incorporate existing regulatory requirements. One of the things that we need to think about is how do we provide guidance to organizations on how to better merge and meld these things so that they can take one step to meet multiple requirements and multiple standards.

Mr. Stine added two points:

- I think the point Adam made on information sharing is right. When you look at the tiers that are currently included within the framework, one of those dimensions across the different tiers is external participation; a narrative describing what you can take from the community and the returning things you've learned to the community. You can have those types of sharing
- We also have a separate body of work that focuses on improving integrated cybersecurity and broader enterprise risk management discussion (ERM) which is looking not just at cybersecurity, but cybersecurity alongside reputation, financial, operational, etc.… There are considerations around your organization, your enterprise's role in the broader ecosystem, considering, for example, risk effects on your organization and how they can have other effects from a national perspective. That could be an area that we can relate more, bring some of those concepts into the CSF, or at least point to these other resources as well.

Mr. Groman asked if that is currently in the ERM?

Mr. Stine replied: We do have some of those considerations in there, but maybe we need to call them out more explicitly.

The Chair compared the RMF and CSF: Regarding the RMF, my impression has been that the RMF is directive in terms of making risk management decisions. If your information has some level of sensitivity (low, moderate, or high) then that directs you to have a program that implements certain controls. I think of the CSF more as starting from scratch. What is the sensitivity of your information? What do you think the risks and threats to that information are? What controls and processes do you need to implement to mitigate those risks at an appropriate level? My impression is that the CSF is much more flexible to different users, different environments, and different management decision making. I think that's part of why industry has liked it and adopted it. It may or it may not be what the government or what the Congress intended with FISMA, but I've always seen that as a little bit of a disconnect.

Mr. Sedgewick responded: I don't want to speak for the RMF team, but I would agree with your description of the two. I certainly agree with your description of CSF. I think that's what people love and hate about it. Because it is so flexible. I think there are ways that we can work with the RMF and the CSF, so you get the benefits of both. There are ways we can bring those into better alignment so it's clear how to use both together. I think that's the discussion we'd like to have. Can we benefit from stakeholders that are using both? How have they done that internally? Can they share their lessons learned?

Ms. Pascoe added: Federal agencies will be the key stakeholders for that. I don't know how many organizations outside the federal government environment would find that useful, but we need to try to drill into that as we move forward.

The Chair indicated it would be useful for the board to hear more on that RMF-CSF interaction and dynamic moving forward.

Mr. Groman added: My question was prompted by the questions you are receiving, not by the document. I'm an enormous fan of CFF, CSF and the RMF and the work that you're doing at NIST. When people come to these documents and want it to be easier and simpler and wonder where's the answer? It's not going to be there. That's not your role to lay it out. The decision of how to implement is going to depend on a lot of factors. It won't be the same for the CIA and the FTC, or for a different company. But the part about the risk needs some work because this contemplates that different entities will have different willingness to accept different levels of risk. One organization may have a higher risk tolerance than another in the exact same area. they'll have a different risk tolerance and if their tolerance is higher, they may choose to have less cybersecurity. That won't take into account the other issues, externalities, or other stakeholders. I don't know where that gets woven in but, as we move forward as a country, that's going to become more and more important.

Ms. Pascoe replied: This is something we're considering. I'm also on the team for the AI Risk Management Framework. This is an important consideration for that framework as well since different departments and organizations have different risk tolerances. The board may have a different risk tolerance than other IT professionals. Your regulators will obviously have different risk tolerances and thresholds that they will set and then your customers will have different expectations. The individuals and broader society could also be affected. We did get a little feedback from organizations asking what it means when you say the CSF is risk based? What does that mean? How do I determine my own risks? That is an area that we're going to focus on, and part of the answer may be the RMF but that's not the whole answer.

Ms. Fanti asked for clarification on alignment: Could you comment a little bit more on the issue of alignment? Concretely, what are you supposed to be aligned with and what kinds of changes would that result in?

Ms. Pascoe replied: We have not yet decided exactly what alignment means. It may mean something different for each topic that we're addressing. The way that we're going to be organizing internally is for each of the themes that we're discussing today, you'll know that the themes align with the NIST cybersecurity program. Everybody that has a major cybersecurity publication or framework at NIST will be part of the process to update the framework. This is a holistic effort across NIST, to work together to figure out what updates we want to make it more aligned with some of our other work. We will likely also have to ask stakeholders what they consider alignment. Does that mean that information from all the resources are put into the framework? Does that mean that we draw better connections between frameworks to make them more interoperable or so that stakeholders know how the frameworks are supposed to be used together? For each of the themes, this may vary in the extent that we change the CSF core versus the CSF narrative versus the resources associated with the CSF. For example: we've gotten a lot of new guidance under the executive order on supply chain software security. Should that be incorporated into the framework? Do we need to figure out how to map and make this easier for folks to use together? That's what we're going to be focused on and it's going to be a hard process to figure out how to do that.

Mr. Sedgewick mentioned: We have new online informative references that might help organizations understand how to use these things together.

(Back to slides)

- Theme 4: Ensure the CSF remains technology neutral but allows it to be readily applied to different technology issues – including new advances and practices.
  - Making sure that the framework addresses IT, OT, and IoT secure software development
  - Given guidance from Executive Order 14028, commenters suggested that additional guidance be given to software security.
- Theme 5: Emphasize the importance of measurement, metrics, and evaluation in using the CSF
  - There were a lot of comments about measurement and assessment of the framework and cybersecurity programs more generally.
  - Sub-theme 5.3 covered expanding, or removing, the tiers, and consider it a maturity model or don't do that. What to do with the tiers is a big topic and how we address the topic of cybersecurity risk assessment and measurement more broadly.
- Theme 6: Consider cybersecurity risks in supply chain in the CSF
  - We heard a lot about what's on supply chain cybersecurity, and how to address third party risks within the framework.
- Theme 7: Use the National Initiative for Improving Cybersecurity in Supply Chain (NIICS) to align practices and provide effective practices, guidance, and tools to bolster cybersecurity supply chain risk management.
- Our Process and Next Steps
  - Stakeholder driven approach.
  - We will have workshops, virtual and in person, roundtables, and one-on-one meetings to try to engage stakeholders and fill in gaps in the comments.
  - Stakeholders will drive the extent of the changes.
  - We're trying to be targeted in our outreach.
    i. We're going to try to figure out how we can do in-person workshops, but I think we're in a new reality where we rely on virtual workshops and, maybe, Adam or I fly out to somebody else's conference, and we engage in that sector or that issue,
    ii. We will not have six or eight in-person workshops, but we hope the stakeholder engagement will still be high,
    iii. We hope ISPAB can help us with getting some feedback and having good conversations about the update.
  - We will have a virtual workshop soon.

- We plan to have at least one draft for public comment:
    i. Before we issue that draft, we will likely have something like a concept paper that will discuss each of these themes, the challenges associated with each theme, and some suggested approaches for how we could address each theme,
    ii. We'll take those papers to stakeholders' workshops and see if we can get some feedback on that to build out on the changes that we want to make to the framework.

- NIST CSF Resources
    - We have been churning out a lot of material connected to the cybersecurity framework.
    - Recently published:
- Ransomware Risk Management (NISTIR 8374),
    i. Ties fundamentals found in the Cybersecurity Framework and makes them relevant to the use case of ransomware,
- CSF Profile: White House Fact Sheet (by Seamless Transition),
    i. Some are developed internally and some of these things come from external stakeholders who have a use case.
- CSF International
    - We did get a lot of international feedback. Most of the comments were to continue or elevate the work that you do internationally and have international discussions.
    - We have quite a few translations and adaptations.
    - International engagement is reflected on our website.
        - ✓ Historically a lot of the translations we receive are developed by external stakeholders. A great example of that is our Ukrainian translation, which we received, and we just do a quality check and put it on our website.
        - ✓ The State Department is funding translations and same with the International Trade Administration within Commerce. So some of the translations, not just the Cybersecurity Framework, also the ransomware guide, privacy framework, a lot of our Quick-start guides are getting funded by the State Department.
        - ✓ Expanding this partnership to also include the engagement webinar and direct conversation side and in multilateral venues such as standard SDOs as well as places like G7, G20, and OECD

The Chair suggested: To the extent possible, I think detailed case studies, actual employee application by organizations, are probably very useful for something like the Cybersecurity Framework. Who's used it successfully? What did they think about that you didn't think of? That's probably very helpful, just as a resource.

Mr. Sedgewick replied: I think that explains a lot of the appetite for the workshops, because that's what a lot of people were getting at in the discussions. Being able to connect with your peers and comparing notes. We have some of those. Some are short but that is something we are interested in expanding.

Ms. Pascoe added: We call them success stories on our website. We received a lot of comments from associations, but we received a lot of comments from individual organizations to talk about how they use the framework in their own company, and the benefits and challenges that they face associated with that.

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

Security of Hardware

Sanjay "Jay" Rekhi, NIST

Introduction

Jay Rekhi

- I've been here almost one year
- Fairly new to the federal government, but lots of experience in the retail industry and before that in Cypress Semiconductor.
- Role is to combine all that experience and do hardware security.
- I would like to get your input as we put this program together.

Hardware Cybersecurity Program

- Semiconductor is pervasive for critical commercial and military applications.
- Hardware stays with us for a very long time.
- The vulnerabilities are hard to find, and even harder to mitigate. Not easy to put a patch on a hardware system. There are ways but not all the time.
- Our objective is to put together some of the research programs, put the documentary standards or the best practices, the reference guidelines, and show the world how to do it.
- The program entails figuring out, with our industry partners, some of those best practices for trust, security risks, and vulnerability management across semiconductor development industry so we can start putting them together.
- Establish the threat models and figure out how you would mitigate them.
- You can put some of that knowledge into tools. There's been a huge progress in combination with DARPA and DoD in some of those areas.
- A lot of the companies are responsible for IP protection. How do we protect the IP when the things are getting designed all over the world?
- I'm amazed at the scale and the speed at which some of our programs run.
- We've already demonstrated some success in this area and I'm hoping to be able to repeat that with a slight adjustment
- Some of the cases that we've already started looking at with Intel and companies like Broadcom.
- The DoD is a big supporter of this as well because of our interaction with DARPA.

Problem Statement & Scope

- Semiconductor development is a complex world. It has become more complex as things are starting to scale.
  - Things get designed all over the world. They get assembled all over the world. And by the time you get a chip on your board, it has already been around the world at least once.
  - It's a problem, not only in the development, but also in manufacturing.
- The Problem:
  - Securing the device itself,
  - Security of the semiconductor,
  - Vulnerability management,
  - Making sure that communication on chip and off chip is secure. That you've got the right protocols in place. You've got the right encryption, etc. in place.
  - Assembly. As I assemble the IPs from various folks, this is going to get even more complex. We start to put things together in complex packages. There may be more chances of people contributing.
  - As the world's supply of semiconductors dwindles, this risk continues to increase. Somebody could put something in there that may steal information.
  - Supply chain component:
    - ✓ Part of the value chain of semiconductor,
    - ✓ Manufacturing typically happens through contract manufacturing,

- ✓ Today U.S. does less than 12% of overall manufacturing of semiconductor. Almost 60% of that happens in Asia: China, Taiwan, and Korea. And 0% of the advanced fab happen in the U.S. The U.S. is about two or three generations behind the most advanced fab owned by TSMC. US folks have started manufacturing in other countries so that there is a huge risk in the security of the materials that gets supplied. The stuff that gets put on the chip, either intended or by accident.

- We also would like to secure the manufacturing of the semiconductor.
  - ✓ The fabs used to think they were better controlled than Fort Knox. Nobody could get in, but they forgot everything is still connected over the internet.
  - ✓ There have been a few cases where the fabs have been affected.
  - ✓ TSMC is putting together a process using the CSF. Hoping to be a live case for that in future meetings.
- Security for this is a complex problem. It starts with inception, and it goes all the way to the end of life and everything in the middle.

Ms. Fanti asked about manufacturing risks: Do you see the risks in the manufacturing column different from those analogous risks in other IT systems? Is there something about the fact that it's for hardware that that makes them different or more challenging?

Mr. Rekhi replied: Most of the IT security risks are about the same, with some specialized equipment that only a few companies use and there's special software. It's very custom so they make their own patches. And how the patch is implemented, that's different again. They use the same I.P. so the same risks show up over there. That may be additional risks that we'll have to look at to address your question.

Creating Helpful Incentives to Produce Semiconductors for America (CHIPS Act)

- Supply chain issues:
  - There's increased interest in silicon and security of supply chain because of the reduction in number of chips that we get. It's a $52 billion program over 5 years. Intention is to bring manufacturing back to the U.S. It amends the National Defense Authorization Act (NDAA) and includes money for three major programs:
    - ✓ Financial incentives:
      - i. For companies to make their fabs in U.S. or modernize their existing fabs. It is a $39 billion program,
    - ✓ R&D program: $11 billion,
      - i. The National Semiconductor Technology Corporation will be a fab house that's going to be built purely for R&D.
      - ii. The advanced academic section: They want to do research into what some of these advanced technologies are where, instead of trying to scale silicon you can put that on a package.
        - Improving the capabilities of manufacturing in the USA.
        - Improving the science of measurement so we can continue to advance and put together some of those techniques to make semiconductor more economical and viable.
        - NIST will play a very big role in this.
    - ✓ Workforce Development
      - i. Workforce is a major component because the jobs don't exist here, so those have to be pulled back.
  - The last $2 billion in there is for DoD.
    - i. DoD is expected to secure the chips and commercialize the R&D that goes behind it

ii. I'm hoping, because of the security component, we get to partner with the DoD on some of these aspects.
iii. Because of the partnership with DARPA, I'm hoping some of that capability could be demonstrated very quickly and we get buy in from the Department of Defense.

- Collectively, the objective of these programs is to improve and bring some of that manufacturing back to U.S. and security and trust outcomes.

Ms. Moussouris question on environmental impact: Are there any funds specifically allocated towards looking at the environmental impact of manufacturing domestically and potentially creating incentives for deliberately green and sustainable technology in US based manufacturing? I know that a huge problem with the existing manufacturing plants around the world is it has an incredibly negative environmental impact.

Mr. Rekhi replied: I don't really know the answer to that question. I'll have to get back to you on that one. I'll look up the act. That's a great question.

NIST R&D for CHIPS

- NIST has a central role in all of this.
- Significant funding for this program.
- Established a task force with NIST building the blueprint on how this will be executed once the Act has passed.
- Connecting with the research facilities, both inside and outside of NIST to assess the current state and what's needed.
- Establishing connections with various stakeholders with the focus on Carnegie Mellon, the University of Florida, Intel, Qualcomm, DARPA, and some of the government agencies, who are really interested in working with us and trying to improve the security and trust in semiconductor.

NIST Workshop on CHIPS

- In April all the labs came together in a workshop.
  - It was done over four days and we ran eight different panels.
  - These panels were divided in two sessions:
    - ✓ The first session was a presentation from the panelists and some discussion around that, and
    - ✓ The second session was open to the public where we collected their inputs and priorities.
  - One panel was focused on Security and Trust in Semiconductor:
    - ✓ Brought in experts from various fields,
    - ✓ The first session was moderated by MITRE,
    - ✓ Great conversation on determining priorities and who's responsible for security,
    - ✓ Got valuable input from all our stakeholders,
    - ✓ We've come up with some grand challenges across NIST. One of them is focused on security and trust. The second one is going to be on the data sharing aspect.

Security and Provenance of Semiconductors

- Challenge: To advance the state of the metrology so we can enhance the security, trust and provenance as the design moves across the platforms.
- The design and development are going to be diverse, so the capability that we need to build is also going to be diverse and very comprehensive and needs to be formalized.
- Going back to what are the threat models that various industries have, what capabilities exist, how do we build them in?
- Design tools that can build security in, so you don't have to think about it.
- Some of the ideas from SSDF because a lot of the design that happens today is software based.

Path Forward

- Identify some of these formal methods, for example: threats
- The methods should improve the data analytics and automation from the inception of semiconductor all the way to the end of life.
- Identify what are some of the trusted technologies that we could use. Is AI and machine learning something that could be leveraged? Those are some of the things we will be researching and I'm willing to partner with a lot of folks on coming up with some of these trusted technologies that may already be out there.
- Build and document standards.
  - We're hoping to leverage a couple of things that are already happening within the DoD. Either build the standards with them or contribute to the standards that are already being built.
- Grand challenge: Standardizing New Materials, Processes, and Equipment for Microelectronics
  - Our problem is that, in the semiconductor world, the data is really siloed. Even within the same company, the data is not shared. This means that security issues can go unnoticed until you're very late in the process. It also means that there may be some innovation that could have happened based on the data and analytics that you could have done while the process was happening.
  - We've been asked to build capabilities in this area.
  - As a designer back in 2007, I never knew what was happening in the fab. It was only after we got the first wafer that I could see and figure out how close or far away it was from what I designed. I think the problem still exists and it's even more complex. We know there are a couple of things that exist out there that we could use.

Summary

- I described:
  - The hardware security program, what it is trying to do,
  - What the US government is asking folks to do. and
  - Where they're trying to put their money.
  - How NIST is involved.
- The program that we started to build together is very similar to what we've heard from our stakeholders.
- I mentioned a few research ideas that I will be pursuing with various folks.
- As a next step, I will be building business cases and align them to the language of the bill.
- When the funding starts rolling in, we can accelerate many of these programs with our partners. Everybody's itching to go forward. There's a lot of capabilities out there. It's about time we really put them all together to good use and improve the state of security and trust for hardware.
- I'm open to feedback. What is it I haven't thought of that I should?

Discussion

Ms. Moussouris asked about data sharing: It seems like the reason for the lack of data sharing is due, in part, to privacy and concerns about leaking trade secrets. What do you think is the likelihood of adopting technologies that protect this data with cryptographic privacy guarantees compared to statistical guarantees? Is there any hope of anything less than cryptographic protection and what would that mean in the context of sharing within a company for example?

Mr. Rekhi replied: I don't really have an answer to that question. I've seen both, use of cryptography but you need to have an NDA. The statistical one is a relatively new area that I'm starting to see. In fact, there's a professor in University of Texas at Austin who was doing something with Samsung to improve all the masks, because some of these things are happening at such fine geometry that even a small change in mask may just blow up the whole chain. So, he's using some of the statistical techniques. I'm hoping

that would be a good case as more folks see how some of this data sharing can improve the yield of the systems because it's got a dollar value attached to it.

The Chair asked two questions:

- **Side channels**
  - One of the things that came up in some meetings a few years ago was this whole matter of the side channels in semiconductors processing sensitive event information: Specter and Meltdown vulnerabilities. Seems like that might fall into the security engineering, SSDF, etc.
  - Two questions:
    - ✓ How serious are those kinds of problems taken by semiconductor manufacturers today?
    - ✓ Are those problems a matter of applying what's known and putting more effort into the problem or is research needed to get to the point where you can reduce or mitigate those sorts of channels being introduced?

Mr. Rekhi replied: That's a very loaded question. I'm someone trying to address them piece by piece. They're doing something serious so that they can catch and mitigate risks like Spectre and Meltdown through some software designs. And hence, you're looking at, once the chip is out, how can you continue to try to mitigate the vulnerabilities that you'll find there? This is part of the thought process right now. Do we have solutions yet? No. I know, there's some research happening in George Washington and Carnegie Mellon to try and address some of these issues.

In response to the second question: as more and more secrets go onto the chip; folks are looking at the techniques on how they can obfuscate some of them such as Physically Unclonable Functions (PUFs) and use all these obfuscation technologies. One of the things that I noticed very recently is a tool that somebody has designed where, once you identify what's critical, they can try to hide it through different metal techniques or some layout techniques. And then there are areas that people look at heat to see where your key may be stored, so folks are starting to put heat sources so they can confuse the big adversaries. I think DARPA has done some work in this and they've identified four different risk areas and how to address them. There's a lot of R&D that's happening. Not a whole lot of that capability has gone into the tools yet. We are early in that stage.

The Chair commented: I thought I heard you say that a major approach was just trying to mitigate the Spectre/Meltdown kinds of problems with software modifications once they're discovered in the field. I understand that if you've got some chips that have that vulnerability in them then all you can do is program around it. The reason software vendors don't like to do that is because it may not work, or it may impose a performance hit. It sounds like you're saying that's a lot of research for methodologies to keep those things from being introduced**.**

Mr. Rekhi responded: We are in discussions with our partners to try and document these things and formalize the research.

Ms. Moussouris asked two questions:

- **Supply chain assurance -** I was wondering if any of this is covering supply chain assurance in the sense of looking at cryptographic signing of firmware or even looking into more end to end of broader hardware stacks?
- **Vulnerability Management -** I didn't see vulnerability management addressed. You've talked about updates, but what about vulnerability disclosure and the ability for reporting novel vulnerabilities and having the process to address them in hardware?

Mr. Rekhi replied: On the first question: Yes, we have a person over here who's implementing Root of Trust across the supply chain. This conversation is happening as we speak between Applied Security Division, and our analog division to the communication technology on how that could be used in the non-digital world as well.

On the second question: I think what you're asking is, is there something that holds the semiconductor developers to disclose vulnerabilities? And I think that was one of the discussions that we had in our panel, that not everybody is disclosing that voluntarily.

Ms Moussouris clarified: It's a little bit different. It's the process of receiving vulnerability reports from outside of your organization, analyzing them for security impact, and then acting and providing remediation. That whole process is, sort of, if you see something, say something, but for bugs. I'm just wondering if there's any "building in" of the capability or the awareness that this is something that hardware manufacturers are going to need to adopt, much like software manufacturers are growing in their adoption of this process.

Mr. Rekhi replied: If they haven't already, they should. That's a good input and I'm taking note of that.

The Chair added: In the software world, one of the things that we worry about a lot is third party components of upstream supply chain, in some cases, open-source suppliers. My impression, just from what little I learned about hardware design and semiconductor design in my last job is that you get some of that in the hardware world too. One developer creating a design and then some integrator consuming that upstream and integrating it, so it goes on to the chip. It seemed like that raised all the problems that we have in the software world: did the person who designed it think about security? If they did, did they get it correct? Was it modified as it went downstream to the integrator who ships it? Are people worrying about that? Is there anything you can do if you get this thing designed for a blob of silicon that you're expected to put in the middle of your chip that might have somebody's intellectual property that he didn't want to talk to you about?

Mr. Rekhi replied: Actually, that is a risk and that's one of the big portions of the supply chain where something could be left, either accidentally because of missed bugs or a bad design, or intentionally. So yes. That is a big worry. Are we doing something about it? There's a lot of research that's happening in the various parts of the world in trying to address that. There are some methods that one of the labs in NIST has come up with. Can that be deployed across? We don't know yet. Intel is going to be publishing a paper on how to address some of that. Do we have a comprehensive approach so that this doesn't happen? I think the answer to that is no. But it's something that we are very, very interested in.

Ms. Fanti asked about the use of machine learning: Some of the technologies that you've been mentioning as countermeasures, for example, PUFs and machine learning, are tools that have been used in cat and mouse races for defending systems and then an adversary that gets access to that system is typically able to generate inputs that can fool the system. So first, I was hoping to hear your reaction to that. But also, I wanted to point out that there has been a trend, at least in the machine learning community, of trying to develop certifiably robust machine learning models that are provably robust against an entire class of attacks. If you're outside the class of attacks, you have no guarantees, but at least within that class of attacks, you can be sure that there won't be any inputs that can trick or fool your classifier. I was curious if that style of research is being encouraged in this new research initiative you're putting forward?

Mr. Rekhi replied: I know there are folks on my team who are looking at how they can build those models and learning capabilities.

Ms. Fanti asked about supply chain root of trust: You mentioned about the supply chain root of trust. I think that's a brilliant idea. One of the things we had heard when standards for such things were being written was that it's very expensive to deploy. Is there going to be a feedback loop in your research about how implementable these ideas could be, beyond the security value added but whether they're feasible for the industry? Doing supply chain with root of trust for measurement adds a lot of overhead to your manufacturing process to be checking devices in and out and checking that hardware root of trust as you go, but I feel like here is a great place to spend that money.

Mr. Rekhi replied: Synopsis is doing work in trying to make security more pervasive. I don't know what type of problems it solves yet. A lot of research is needed to make it more cost effective. And who can afford it? These chip manufacturers, and I know that Cypress Semiconductor, the company I used to work for can't afford some of this. They make chips that you sell for less than a cent. Those are some of the chips that are most vulnerable. A one cent chip was a clock chip between everybody in the world, I could easily insert something in there. I'm hoping some of the regulatory bodies are coming up with some standards. There is a conference happening at the end of this month, sponsored by DoD, to help put together some of those standards.

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

AI Security Metrics and Threats

Harold Booth, NIST

Disclaimer

- This was originally a talk I gave at RSA.
- How we measure things is important to think about when looking at security. We need to understand:
  - What it's a measure of? What those units are? These are the kinds of questions we need to be asking ourselves in the physical and in the virtual world. For example:
    - ✓ Do know what a Newton is? Newton is a measure of force. It's mass times acceleration, so it's kilograms times meters per second squared. Systems that are described by those may be slightly different. 10 Newtons may be 10 kilograms accelerating at one meter per second squared or it could be one kilogram accelerating at 10 meters per second squared.
    - ✓ This also matters in virtual world. There's a need to really understand what the metrics mean when we're looking in the AI security realm

Acknowledgements

- This work couldn't have been done without contributions from the rest of the team. I just want to make sure I acknowledge that. We also used some publicly available datasets sources.

Introduction

- Will be covering adversarial machine learning briefly.
- Quickly touch on Dioptra.
- Discuss a couple of use cases that dive in on some of these ideas of measurement, risk, and that sort of thing.

Adversarial Attacks (Draft NISTIR 8269)

- NIST published Draft NISTIR 8269 in 2018 and included a taxonomy and terminology diagram for adversarial attacks and defenses of machine learning. Below, see Figure 1. Taxonomy of Attacks, Defenses, and Consequences in Adversarial Machine of (Draft) NISTIR 8269.

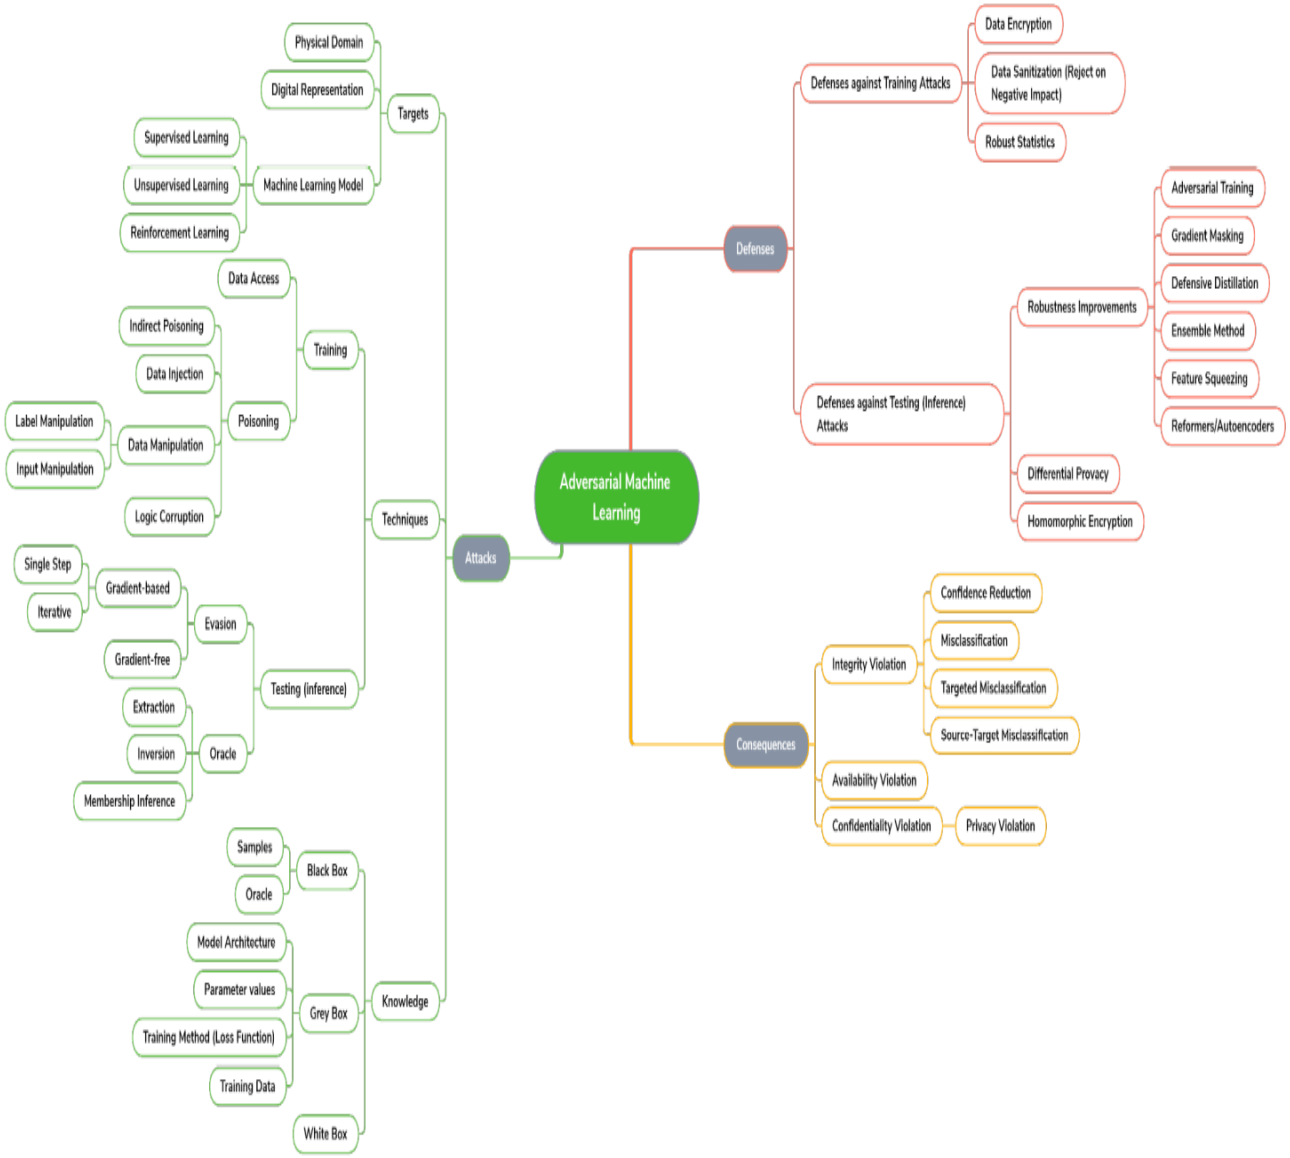

Below, see Figure 2. Taxonomy of Attacks, Defenses, and Consequences in Adversarial Machine of (Draft) NISTIR 8269

1. Attacks
    a. Targets
        i. Physical Domain (of input sensors or output actions)
        ii. Digital Representation
        iii. Machine Learning Model
            1. Supervised Learning
            2. Unsupervised Learning
            3. Reinforcement Learning
    b. Techniques
        i. Training
            1. Data Access
            2. Poisoning
                a. Indirect Poisoning
                b. Direct Poisoning
                    i. Data Injection
                    ii. Data Manipulation
                        1. Label Manipulation
                        2. Input Manipulation
                    iii. Logic Corruption
        ii. Testing (Inference)
            1. Evasion
                a. Gradient-based
                    i. Single Step
                    ii. Iterative
                b. Gradient-free
                c. Oracle
                    i. Extraction
                    ii. Inversion
                    iii. Membership Inference
    c. Knowledge
        i. Black Box
            1. Samples
            2. Oracle
        ii. Gray Box
            1. Model Architecture
            2. Parameters Values
            3. Training Method (Loss Function)
            4. Training Data
        iii. White Box
2. Defenses
    a. Defenses Against Training Attacks
        i. Data Encryption
        ii. Data Sanitization (Reject on Negative Impact)
        iii. Robust Statistics
    b. Defenses Against Testing (Inference) Attacks

i. Robustness Improvements
    1. Adversarial Training
    2. Gradient Masking
    3. Defensive Distillation
    4. Ensemble Method
    5. Feature Squeezing
    6. Reformers/Autoencoders
ii. Differential Privacy
iii. Homomorphic Encryption

3. Consequences
    a. Integrity Violation
        i. Confidence Reduction
        ii. Misclassification
        iii. Targeted Misclassification
        iv. Source-Target Misclassification
    b. Availability Violation
    c. Confidentiality Violation
        i. Privacy Violation

- Techniques include Poisoning, Evasion, and Oracle attacks.
    - Various kinds of techniques when attacking in terms of machine learning.
        - ✓ **Poisoning:** Where data is put into the training set that can later be used by the adversary to cause a classification that's beneficial for the adversary.
        - ✓ **Evasion attacks:** An adversary changes the input or perturbs the input in some way that would make it, so the classification happens in a way that's favorable to the adversary in some way.
        - ✓ **Oracle:** attacks are when you're trying to reverse engineer exactly how the model was trained.
            i. Membership inference attacks can identify whether a particular individual or row was used within the training set.
            ii. Model inversion where you try and rebuild the model in some way and try to understand what the training set is or how that model is performing.
    - Things you can do to try and defend:
        - ✓ **Data pre-processing:** These are things that would help with evasion attacks, where you can do sweeping and some other techniques around that.
        - ✓ **Data Sensitization:** For data poisoning attacks
        - ✓ **Provenance signing:** In the traditional IT space, you could do provenance signing of the training data and make sure that the data is what you think.

- Consequences include Confidentiality Violation, Integrity Violation, and Availability Violation.
    - **Confidentiality:** we can have issues with confidentiality to the model or the training data that was used to create that model,
    - **Integrity:** integrity issues with confidence in how that model is going to perform versus misclassifications.
- **Visual Sampling of Attacks:** A lot of the research is in the visual modality.

- **Patch evasion attack:** A small patch is put onto an image that can cause misclassification. Slide shows an image of a beagle dog that's being misclassified as a toaster.
- **Noise evasion attack**: Lots of pixels are changed on the image, sometimes in a way that is not visible to the human eye, that can force misclassification. Slide shows two images of a car, one with noise added. You can see that it's been it's been perturbed, and this one means a sports car or purse, depending on whether you're protected or not.
- **Poisoning attacks:** Types of samples are put into the training that would be beneficial when you're trying to do the inferencing. The slide shows the target image of a currant being "poisoned" by an image of an apple. In this case, maybe using steganographic techniques, I combine an apple with a currant. When presenting the model with the currant, the inference will result in an apple
- **Oracle attacks:** Where we're looking at things like tabular data used for training the models.
- Visual Sampling of Defenses
  - ✓ **Image pre-processing:** We remove a lot of extraneous features and focus on the salient features that we think would be most useful when trying to get inferencing.
  - ✓ **Desensitization:** We're taking an image and trying to remove extraneous marks that could be backdoor attacks where pixels are there that would be recognized by the models.

- Evaluation Metrics
  - There are lots of metrics, we're looking at in AI. Metrics include:
    - ✓ Accuracy,
    - ✓ Effectiveness,
    - ✓ Sensitivity,
    - ✓ Transferability,
    - ✓ Generalizability,
    - ✓ Perturbation Distance,
    - ✓ Time and Resources.
  - Many see accuracy as the main one but combining accuracy with the others is very important, , especially when you're concerned about security and the kinds of attacks that you could be vulnerable to.
  - The basic take home message is there's lots of metrics and measurements out there. Really give some thought to what you're using and make sure you understand how those play with your deployment.

Ms. Moussouris asked for clarification on content of the slide: Just on that last slide, can you just verbally talk us through with what's in yellow because I can't see those images.

Mr. Booth replied: The idea was that one attack could be useful against multiple defenses. There's a cost when you're setting up a particular attack. One of the things that may be interesting is to understand whether a given attack is useful across multiple models. The three swords with a shield depict the idea that you have one defense that is good against multiple types of attacks. Often, some of the defenses that you have to employ might result in some form of accuracy loss. There's a trade off in some cases. It's not universal.

- Details on Metrics
  - (Slide shows a graph with the "circle of robustness")
  - This graph shows a prediction on the number of perturbations that are possible before I can get misclassified using the dog image. One measurement you can think about is how many perturbations are possible before I can be forced to be misclassified. The types of evasion attacks may result in different forms or different measurements.
  - Looking at perturbations distances, for example:
    - ✓ When you're looking at things like Manhattan distance, which is summing up all the differences across all the various dimensions,

  - ✓ Euclidean distances, like the difference between the original and doing a straight line.
  - ✓ Infinity is looking at the maximum magnitude change across all the changes that you made for a particular image.
  - ✓ If I made a big change to a particular pixel that changed the RGB value, that would be result in much higher L infinity than if I made lots of small changes.
  - ✓ The basic idea when thinking about metrics is combining them.
- Deployment Context
  - What is the text that you are going to be using this for?
  - Can you control the environment?
  - Is this an open environment or a closed environment?
  - What are the adversaries' objectives? What are they trying to subvert in terms of your system?
  - What components does the AI depend on?
  - What are the upstream components? Conversely, what are the downstream components, like when you're trying to make a prediction with an AI system?
  - Think through and understanding your risk profile
- Working on Dioptra
  - Dioptra is a test framework for doing evaluations of AI systems. We built this to simulate attacks and defenses and be able to do training models and then switch them out in various ways. We set things up so we can have multiple training architectures, different data sets, different kinds of attacks, and different kinds of data augmentation.
  - We couple that with different kinds of inference pre-processing, and we mix and match and change all these things in some sort of combinatorial fashion. Combine this with different forms of metrics that would be useful in trying to identify whether a particular model, architect or defense was good or bad or whatever. The intent of this tool was to make it easier to run through those combinations
  - We've been looking at this primarily from a research perspective. Looking from our own internal risk perspective, as we run through evaluations. And from first- or second-party testing.
    - ✓ If I'm looking to purchase a particular model or use a model, I can use this to run through a bunch of different kinds of tests.
    - ✓ If I'm a first party, we're doing development and I can run this as part of my DevOps pipeline
- Use cases:
  - Secondary testing: We're imagining that we are a CTO at a supermarket chain, and we are trying put in a place a new AI system that will allow us to do pricing of fruits and vegetables at a checkout. The Task:
    - ✓ Instead of having to type in the number or ask the attendant, you could have a camera take a picture and it would identify it as a fruit and it would look up the price based upon the fruit and then based on the weight on the scale, calculate the total price and move on.
    - ✓ We identified a high level recognition process as the task.
    - ✓ Is AI necessary? Many times people are quick to use AI when it's perhaps not appropriate. The first thing that I think everyone needs to ask themselves is do I need to use AI? For the purposes of the demo, the answer is yes.
    - ✓ Identify threat and deployment assumptions. Identify which attacks are relevant based upon those threat and deployment assumptions. Then identify what metrics are applicable to those highest priority risks. What is important?
  - Last, build the experiments and then do some analysis and get results. In this case, the misclassifications could be costly if for example, I have grapes and I priced them as cherries. That's bad for the for the consumer but good for me. If I have cherries and I priced them as grapes, that's bad for me, good for the consumer. We don't want either of those.
    - ✓ Threat/Deployment Model: We have an attended self-checkout. We're going to have a human there that's keeping an eye on things in some way. No online learning, we are not going to try

and learn as we go although there is the whole issue with drift. For example, seasons are going to be different depending on what fruits there are. So that may be something they may be concerned with at some point.

- ✓ Attack Profile: In this case, we also are going to have freshly generated digital representations so digital manipulation techniques are difficult to do in this deployment setting. We are also going to assume the training data itself is not particularly sensitive.

Mr. Groman asked: What is the term you're using for the output or the intent or purpose of the AI

Mr. Booth replied: I've been using inferencing or prediction generally. Is that what you're asking?

Mr. Groman clarified: So in your example, the purpose of the AI is to…

Mr. Booth replied: In this case, this is taking an image of a fruit or vegetable, and then it's predicting which one it is. So it's identifying whether it's an apple or grape or whatever.

Mr. Groman: We would call that the objective.

Mr. Booth: This is the task that the AI system is trying to do. So purpose could be, I want to make it easier for my customers to get through the checkout quickly. I'm using tasks to be focused on what it is that we're asking the AI system to do.

Mr. Groman: In your case, is the AI deciding or is it facilitating a decision that is made by a human?

Mr. Booth: It is making the decision of what fruit it is.

Mr. Gattoni added: That's the discrete nature of the task. It makes that decision but then it's part of a larger process that assigns that decision to the value the grocery store charges the customer. The discrete model's task is to decide what fruit it is precisely and then complete the transaction.

Mr. Groman: So, to the extent that an error is made in that task, it could cascade through other issues, not just about price, but my inventory is going to be wrong, and I'm going to order wrong.

Mr. Booth clarified: We'll cover the price thing.

Mr. Gattoni added: The controlled environment slide from earlier is important here because that's the attended self-checkout part of the non-machine measures that are in place in the overall process to help prevent or mitigate the potential for the machine making a mistake. One of the things we look at in cyber when we're looking at automating a task is the benefit and regret intersect. That is the mitigating controls put in place for the regret. In this case, we may be concerned about data poisoning. It could be that something was put into the training data that could be a concern. In this case we had this patch attack, where you could force misclassification by having something within the image. We're imagining a case of having a keychain within the field of view of the camera that could cause a trigger with respect to the classification

Mr. Groman asked about data poisoning: Is the idea that you could start off with data that is itself inaccurate? Is that part of the analysis? Or you're assuming accuracy because the threat later would be making inaccurate? Or is the possibility that you're starting out with the wrong data? Is that contemplated in your presentation?

Mr. Booth replied: Data poisoning itself would be the idea that the data is wrong with respect to the training of it.

Mr. Groman: Data poisoning strikes me as something intentional by a threat or threat actor as opposed to, I am doing something significant with data that can have a significant impact but the data I started with was bad.

Mr. Booth: In this demo, you could make the argument that the data was bad. The data we use for this demo should not be used in the real world. After you generate your model, you should be checking to

make sure that it works across many different domains and in different situations. I'm not covering that in this talk. There's lots of other discussion about that outside the security space, so I'm trying to focus on the security aspect, but certainly what you're discussing is a hot topic of conversation that comes up a lot. Some of what you're talking about comes in the space for bias. Being able to trust that a model is behaving appropriately. Making it so these models can tell you when they have no idea what the heck this is, and they're not going to try and give you a prediction. There's lots of issues with respect to when having this very data intensive process and making sure that it's within the distribution of what you're expecting the system able to see. A lot of these errors are when you get an input that's outside the distribution of when you train the model.

Mr. Groman: Say you're training the model and get to a pomegranate and tomato, and it picks the wrong one. You're saying what it ought to do is say, "I'm not sure?"

Mr. Booth: Yes, or you have better data that can help you disambiguate between a pomegranate and tomato. I would argue if it's never seen a pomegranate and it is training data, it should definitely tell you that it doesn't know. That is something that would be good for a model to be able to do. But if pomegranates are in the training data, and it still doesn't know, that means there's something wrong with the distribution of the data that you fed your model and your training of it.

Mr. Groman: That's outside the scope of security analysis?

Mr. Booth: There's a relationship there in that if I have a bad model, then an adversary can more easily attack that.

- Relevant Metrics
  - We have a couple different bundles and we're going to try to do some pricing
  - The most important thing when looking at this is what is the cost and what is the variance between the actual price and the price that I have predicted. That's our metric for analysis of these systems.
- Combinations Tested
  - We use Dioptra: this is the experimental setup. We had two different models.
    - ✓ Vendor 1: a simple model using the data we have
    - ✓ Vendor 2: a defended model that uses some sort of technique to help against certain kinds of attacks. In this case, it's defended against patch effects.
    - ✓ We use the Friut360 dataset.
    - ✓ We use patch attacks
    - ✓ Looked at things like clean edges or accuracy, adversarial accuracy, and the cost per misclassification.
- Dataset
  - A collection of fruits and vegetables
  - Pulled out the 10 most common ones to focus on and keep it narrow.
- Model Comparison – Clean Data
  - Comparing these two models the accuracy is similar between these two models
  - Clean tested on un-adversarial examples.
- Patch Attack
  - If we apply a patch to the model, we see a big difference between these models.
- Model Comparison – Adversarial Data
  - The first model, which looked like it was performing better, performs worse for the adversarial examples.
  - The defended model behaves much better – obviously it's been trained.
- Vendor 2 Details
  - We included images with adversarial patches in the training set
  - Model learns to ignore those patches because we give it the correct label when it sees that patch.

  - Note: in this example, training is not a panacea. The same patches were used for training and testing.
  - In a real world case , there's going to be different patches.
- Purchase Scenario
  - We invented this purchase scenario where you go to grocery store, and you buy five each of these kinds of fruits.
  - Added real-world prices to each type of fruit.
- Cost of Cart
  - We have the average cost of the cart, the true cost, and the cost of a cart if we apply our banana patch where everything is priced as a banana, which was the cheapest fruit in our dataset.
  - The undefended model's cart was $8.14.
  - Our defended model's cart was $24.81. You'll note that the price was still a little bit lower than the actual cost.
  - The supermarket may not want to implement either of these systems because the cost is too high.
  - Working through the math and trying to figure it out is super important before you ever deploy these systems and understanding if your regret of implementing an automated system is high.
- Key Takeaways of Supermarket Scenario
  - Focus on impactful operational outcomes. Monetary impacts in this example are going to be key in determining whether it's worth buying and deploying,
  - Identify failure modes most important for your context,
    - ✓ Which attacks account for most of the risks that you're going to have
    - ✓ Include all aspects of the system, including the humans. Having a human there that's looking at what's going on at the checkout. There's research about having humans together with the system performing.
    - ✓ Maybe not having a fully accurate system coupled with a human might be better than having a fully accurate system. That may be another way to look at this.
  - System context can offer various mitigations.
    - ✓ Technical: integrate with the patch detection model
- Integration and Regression Testing
  - Follows roughly the exact same risk model: Identifying the task. In this case, we're imagining ourselves as the developer that is going through and trying to create a library for doing road sign detection. As a developer, I don't know precisely how it's going to be deployed out in the world, but the task I know I'm doing is detecting road signs: both where they are within that particular image as well as one particular class of road sign. So this may mean that I can be used as part of an automated assistance system like maybe putting a warning light or some other thing to assist the driver or for an automated driving system (ADS) that is making decisions. It would be part of the decision-making process based upon the road signs that my system is detecting that would then decide whether to stop or continue to go through an intersection, etc.
- Threat/Deployment Model:
  - The first deployment model when we're talking about driving systems is the threat of accidents.
  - Second is unmonitored environment.
- Risk that adversaries could do things within the environment. I can put things on stop signs. I can occlude particular signs or cause other mischief. I could have a primary human driver, or it could be backed up depending on if it was for the automated driving vehicles. There are different levels. In levels one and two, you still have the human. Level three, four and five, you start getting into decreasing involvement on the part of the human.
- Attack Profile:
  - We're not terribly concerned with digital manipulation attacks.
  - We're not going to be too worried about training data.
  - As a vendor we may be concerned about model extraction attacks.

  - We are concerned about physical manipulation of the environment.
  - We could be concerned about dataset poisoning like having physical triggers. Example: stop sign that has as a post it note applied to it causing misclassification
- Road Sign Detection Combinations:
  - We had multiple models that we looked at and evaluated.
  - We used the Kaggle roadside protection competition dataset.
  - We looked at different kinds of data augmentation
  - We did inference pre-processing.
  - Looked at a couple of new kinds of attacks.
  - Looked at a couple different metrics.
- Dataset Considerations:
  - Four different kinds of classes:
    - ✓ We had stop signs, crosswalks, speed limits, and traffic lights. You'll note that in the training set, we have an imbalance in the classes. We have a lot of speed limits, but not a lot of stop signs. Stop signs are fairly uniform, whereas the limits are not. Maybe that's not a big deal, but that's something to be looking at. Needed to be careful to divide images appropriately between the training datasets and the test sets or we got too good of a performance because basically the model was starting to memorize, basically those images.
    - ✓ Training set: what you use to train the model.
    - ✓ Test set: is what you use to validate the model's performance. You hold back data that you don't let them see until you are doing your testing. You don't train on it at all.
  - Since we didn't have a lot of data, we did things like flipping the images, changing the contrast and a couple of other little things to increase the amount of data.
- Performance Metrics:
  - The first thing that we did is we ran a bunch of different models and then we pulled out the most important, or the one that seemed to perform the best.
  - The best performing one was this RetinaNet with 3000 training steps. We looked at all these numbers and started thinking about what they mean. Things like intersection over union, mean average precision, and average precision. The point of this slide is to dive into an understanding what do all those things mean? If you don't really understand what those numbers are telling you, you might learn the wrong thing and you may start making the wrong decisions.

Ms. Fanti asked about who is running the evaluation: In your use case, who is running this evaluation? I'm asking this because I don't understand, if it's someone external who's evaluating a model that you've trained, they might not have the trigger data or patch data to evaluate on. Whereas if I'm trying to prove that the model, I've trained is secure, I'm clearly not going to test it on trigger data.

Mr. Booth replied: For this scenario we're imagining we're the developer and we're trying to test our model before we release it to our customers. or before we to basically test it ourselves that it's secure. There's no triggering yet. This is just looking at the result of training my model. I have my test set and I have one model that looked like it was good. Now I'm presented with all these metrics and I'm just saying what do all these metrics mean?

- Robust DPatch Attack: We looked at different patches, not only patches that were trained on our model. Grabbed other patches that we have from some other random model to see if that works. Looked at things like the JPEG compression and spatial smoothing.
- Patch Attack Results: What you're seeing here is I have a patch attack that's being applied. The graph shows no defenses applied, JPEG compression applied, and spatial smoothing. The undefended the model performs better overall when you have no patch attack. The defended models perform worse overall, but they perform better when you have patch attacks. They don't get misclassified.

Ms. Moussouris asked, about the patch attack: In the scenario where you have it mitigated by having human attendance in the supermarket example. Have you studied how small a patch can be that interferes with the results.? I also wanted to know if you had looked at patches that could be physically applied in the moment that may evade even humans who are trained to look for that interference because the patches themselves might have components that aren't quite visible to the human eye or any other things like that. They can be disguised as a clever Chiquita Banana sticker, and you put it on the bananas and suddenly it's a much cheaper vegetable.

Mr. Booth replied: The supermarket example we thought about was on a keychain because you could think people have all kinds of things on their keychain and they would just make sure that it was within the field of view of the camera. We didn't look at trying to print anything out or do anything in the physical real world. I have seen plenty of demos and examples in the academic space where people have done that. Often the criticism is that they're academic. Maybe those aren't real world, practical attacks quite yet. We're looking at this through the prism of these are things that you should be concerned about. Here's a tool that we think that will help you when looking at these kinds of concerns. I personally haven't done an exhaustive study of looking at like what sizes of patches can be most effective versus not the most hidden.

Ms. Moussouris added: That's just something to consider in terms of optical readers that are intended to be used as part of the system as a whole and understanding what spectrum those optical readers can actually intake, as opposed to what the human eye can intake and what kinds of manipulations there may be that are available.

- Backdoor Poisoning Attack:
  - A backdoor poisoning attack can be very effective. The slide shows an image of a stop sign with a post-it note placed under the "TO" in the word "STOP". The image was misclassified as a speed limit sign
- Poisoning Results:
  - The graph shows two models, both poisoned, with regular images and backdoor poison images. For some reason, the model on the right says you're doing well even against the poisoned models
  - Looking deeper, we found that the libraries that underpin these models and how they were generated, had different threshold values for when they would say that something was something versus another
  - The one the left had a 70% threshold before it would tell me that this is a stop sign versus this is a speed limit. The one on the right had a 5% threshold. So that meant that stop signs would appear just about every time and that's why this one looks really good.
  - So if we do an apples-to-apples comparison, now we see that they performed the way we expected where they both perform much worse when I have the poison images that are being provided. We then tried a couple of different models with 35% and 70% confidence thresholds

Mr. Groman asked if changing a confidence threshold in a model is one method for mitigating risk?

Mr. Booth replied: Maybe? I think if you have a higher confidence threshold before you decide, I think you are in some ways reducing your risk. I think I'm comfortable saying that to some degree.

Mr. Groman added: But if you have an algorithm or AI that has a potential impact on an individual, you get a loan or you don't. So if we have a higher confidence level, we'll have fewer, less inaccurate decisions, so fewer people will be wrongly denied.

Ms. Fitzgerald-McKay added: that's one of your risks is that too many people who ought to be approved get denied, right. That's the risk you're accepting by raising the confidence threshold. In your use case for people getting loans, if you raise your threshold then, it makes it a little bit more equal, but the way it might make things equal is that no one gets a loan. Right?

Mr. Booth: In my stop sign example: If I have a high confidence threshold before I'm going to say this is a stop sign, that may mean that I'm going to miss out on stop signs. So that's the concern.

Mr. Groman added: in the case where the more important the decision of an algorithm, the more you would want some higher level of human intervention perhaps. So we're willing to accept or tolerate a higher error rate for getting it wrong if it's a tomato or an orange, but we don't want to accept a high error rate if it's about who gets to cross the border.

Mr. Booth: The point that I want to make is when you have these high stakes things, constructing a system in a way that is intentionally making sure the human is involved and that the system is supporting the human in some reasonable way so that you're constructing and considering all these other risks. I think that's maybe one of things I think that we need to be more comfortable with, thinking about how we construct systems so that humans are engaged, and they are first order entities, and the technology is only a supporting role.

- Key Takeaways from Road Sign Detection:
  - Understand what your metrics are really telling you,
  - Make sure you say how your metrics work and what is being said,
  - Sometimes break them into smaller ones to fully understand what's going on,
  - Understand your tools' default parameters,
  - We got hung up on this because we just assumed that everything was working right,
  - Make sure you know what's going on,
  - That you're doing fair comparisons,
  - Be aware of how dataset characteristics affect metrics,
  - Something to be aware of and be concerned with are class imbalances,
  - Artifacts may exist in your data set. In the data set for the road signs, we had these tracks in images we had to correct for in a reasonable way or else that we would get that information when we're doing analysis of our model.
  - Testing during development can help improve final product. As you're developing, you might identify more and better metrics that help you delve deeper and better into various aspects of the system. You may also identify types of mitigations that can be necessary. Example: if you're concerned about data and discover that your model is vulnerable data poisoning, you may have to go backwards in your supply chain and identify how to better protect that to make sure that you don't have those sorts of issues.
  - Conclusions:
  - There's plenty of things to measure. There's no small set of metrics that are good for all use cases. It's valuable to develop a process for deciding what to measure. I'm focused on those risks when deciding how to test and evaluate systems.
  - You need tools to help you manage all this. There's a lot of complexities, having tools to help manage them is important
- Future Work

Looking forward, we're going to continue to evolve Dioptra. Look at more attacks and defenses. We're looking at more metrics, and we're going to try to increase modalities. I focused largely on doing vision or image recognition. We're doing something with auditory so looking at an audio type task. We're also going to be looking at a tabular data type task. The other thing that we're looking to get to do is expand the community of use. Finding more people that are interested this and helping us and pushing on this and making it better overall. https://github.com/usnistgov/dioptra

Discussion

The Chair commented: The thing I take away from this is my complete ignorance of this area of practice. There's a lot of opportunity for people to use AI naively and get themselves into trouble. If we're

really telling every agency to use more AI to aid its decision making, then it seems there's a there's a potential risk of amplifying the number of people who are being encouraged to use stuff whose implications they may not understand. It's even possible that the organizations that are trying to sell them the tools may not be going out of their way to help them understand. That would really argue for a NIST contribution to help people break that down and understand what questions they ought to be asking. Maybe that's where you're going, but you probably ought to go there faster.

Mr. Groman asked: The notion that in AI, what's taking place is more learning. The algorithm, the artificial intelligence, the model actually has a data set you start with, but it is learning and then creating new actions based on what it's learned, potentially without new data?

Mr. Booth replied: Not always. Not all models are made the same way. There's a difference between static models and models that have online learning.

Mr. Groman: In that context, where your model is potentially learning from data, which means it's creating new data to build upon and then make decisions. That would introduce a whole separate set of risks, or no?

Mr. Booth: Yes, a completely different set of risks. You may have heard of in the news an interesting project, Microsoft Tay, a little AI chatbot that was put out on Twitter. It was supposed to learn how to interact with people through conversations on Twitter. After less than 24 hours that chatbot was posting very racist, xenophobic, antisemitic, etc. messages.

Mr. Groman commented: I think your point is dead on. I'm fascinated by this. The potential downsides are just not fully appreciated and there is so much momentum in both the public and private sector. I can't go to a meeting or a conference where this is not being pushed. I want to do it. I just find that the potential adverse consequences are very rarely addressed in more than a very cursory way and that's concerning. This is just too important for us to get wrong.

Mr. Gattoni added: I think brings into focus Kevin's question about socio impact as part of their assessments and then they gave the early presentation about the AI Risk Framework that's forthcoming. This is something we ought to talk about in our time here and seem to be a good juncture for our skills.

Ms. Fanti added: It might also be useful to think about what kind of risks are happening in the current human run systems and then compare that as well. You're presumably mitigating some risks but introducing new ones.

Public Comment, Summary of Day 1, and Board Discussions

- No public comments were received.
- Board Discussion
  - **Mr. Romine:** As we were getting public input on what the risk management framework for AI needed to look like, we began to recognize that a lot of the comments we were getting were not purely technical. When we put out our initial draft suggesting that we needed to recognize this sociotechnical aspect the feedback we got agreed with us. We're not exactly sure what to do but we'll get there
  - **The Chair mentioned:** I think this was an instance of a meeting where the speakers got some level of feedback from us as we went. Is there anything that we want to say formally, to put on the record? Ms. Moussouris asked for time to think about it and the Chair agreed it could wait until tomorrow.
  - **The Chair asked:** If the board has been briefed on the RMF. Mr. Brewer and Mr. Scholl indicated that they have, but it’s been quite some time and it would be good to bring them back. The Chair agreed.
  - **The Chair suggested** non-technical considerations in the framework or other guidance as a topic ripe for comment.

Day Review and Meeting Recessed

The Chair adjourned the meeting at 4:08 P.M. ET.

Thursday, July 14, 2022

The Chair opened the meeting at 10:05 a.m. ET and welcomed everyone to the call. Started the presentation without a quorum knowing that one would exist prior to any motions would be considered.

Threat Briefing – Current Advanced Persistent Threats (APT) Activities

Daniel dos Santos, Head of Security Research, Forescout

Introduction

- Three parts of this presentation.
  - Data Overview – risk & threats: The data that we observe and where that comes from,
  - Deep dive into APT trends: A deep dive into emerging trends that we have observed from this data and from other sources, and
  - The Conclusion that we can draw from these observations.

1. Data overview – risk & threat

- How we collect data
  - We collect data from our customer networks that we observe
  - Customers that choose to share anonymized data back to us on the cloud,
  - Telemetry data uploads
  - We observe on sandboxes and darknets
  - From our own adversary engagement environment,
    - ✓ A set of honey pots that is deployed on the public Internet and waiting to be attacked and to learn attackers.
- Share that information in the forms of Indicators of Compromise (IoCs), threat feeds, and threat reports to interested researchers, agencies and so on.
- Two main things that we want to discuss are around risk and threats.

Network Risks

- We collect data from more than 18 million devices that are under monitoring.
- We see an average 24% of devices are nontraditional IT devices:
  - operational technology,
  - healthcare devices,
  - internet of things, such as Wi Fi and IP cameras and things like that.
  - It's increasing an already big attack surface that is being targeted by threat actors.
- Risks trends from those devices
  - IT: network infrastructure is among the riskiest devices and that includes things like routers and firewalls. And that's mainly because that is one of the main initial access points for ransomware and older actors nowadays. Network infrastructure is a growing trend among threat actors for initial access points

  - IoT devices: surveillance category - IP cameras, network, video recorders, and the VoIP category are an increasing concern because they have lots of easily exploitable vulnerabilities and often are the most internet accessible devices that we observe on customer networks.
  - Operational technology: There is an increasing concern with building automation devices, for instance: heating, ventilation, air conditioning, physical access control, fire alarm systems, and things like that, that are integrated for control of buildings. They have a critical impact and often have internet connectivity.

APT Origins and Targeted Sectors

- We collect data from malicious DNS requests that are detected on the customer networks, and we see close to 20 APTs and more than 300 malicious domains at any given time,
- The observed APTs in 2022, you will see that all of them are either of Russian or Eastern European origin.
  - One of them is a state sponsored actor,
  - The others are cyber-criminal enterprises targeting financial services, that's banks and other financial institutions or other types of sectors.
- Most attacked sectors
  - Government and financial services are among the most targeted sectors.
  - After that, you have manufacturing technology and utilities (oil and gas and so on)
- APT activity by origin
  - Close to 83% of the activities are Russian or Eastern European
  - After that you have Chinese, Pakistanis, North Korean, and Vietnamese groups

APT Motivations and TTPs

- Two major types of APTs:
  - 47% cybercrime and
  - 53% are groups that are involved in state-sponsored activity
- Types of attacks:
  - Mostly espionage, in some cases destructive
  - Cybercrimes, mostly ransomware, financial scams, business email compromise, and so on.
  - All those groups are still mostly involving Windows-based exploitation, but some of them are moving to target other types of devices such as the ones we mentioned earlier.
- Top ransomware TTPs in 2021.
  - Mostly related to enterprise devices or windows-based devices.

Attack Attempts on Exposed Vulnerable Service

- The attack engagement environment is on the data that we expose on purpose with the goal of collecting attacks.
- We see there are exploits on public facing applications.
  - EternalBlue or Log4J which is six months old.
  - There are classic vulnerabilities that will never go away. Those things are still continually being exploited or adapted to being exploited by threat actors.
- There are also many valid accounts with default, easy to guess credentials for IoT devices.
- Data exfiltration or command and control via HTTP and other application layer protocols.

- Third type of threat actor: Hacktivists
  - Compromise devices to make them part of Botnets for DDoS attacks or to use them as proxies for other types of attacks.
- Attack Origins:
  - The distribution of IP addresses or geography of IP addresses that shows that the majority are originated from U.S. IP addresses.
  - Origin can be difficult to identify due to use of proxies, VPN connections, Tor, exit nodes and so on.
  - After U.S. IP addresses Asian, such as China, Cambodia, and India 30% of the attacks.

2. Deep dive into APT Trends

Ransomware Evolution

- Incidents
  - More than 4000 publicly tracked incidents in the past couple of years related to ransomware.
    - ✓ The U.S. has been a top targeted country with more than 700 publicly tracked incidents.
    - ✓ Non-publicly tracked incidents could be double or more
    - ✓ Currently there are hundreds of active groups, many using the Ransomware as a Service (RaaS) model
      - i. For instance, Lockbit and Conti are among the most active
      - ii. This model is when a group develops ransomware but doesn't actually deploy that ransomware on targets.
      - iii. Relies on a few hits that will penetrate a network and then deploy the ransomware
      - iv. Attackers share part of the of the profit from an attack with the original developers.
- Ransomware timeline from 2019 to 2022:
  - 2019 data encryption
    - ✓ Popular in Germany and is still very popular,
  - 2020 Exfiltrating data before encryption,
  - 2021 Large extortion campaigns
  - 2022 added things like denial of services, public shaming of victims, boasting things on leaked data websites and so on.
- The evolution continues with new sophisticated ransomware families and more types of targets
  - Example: Ransomware outfits work on their own sophisticated encryption implementation software, and they target, for instance, virtualization services, VMware and virtual machines. So not just the workstations directly, but other types of operating systems and other types of assets that are business critical.
  - And we observed that this evolution of the ransomware landscape is far from over and as the adversaries are changing because of at least two things
    - ✓ A proliferation of IoT and OT devices
    - ✓ State sponsored ransomware: a mixture of ransomware with state sponsored activities.

Ransomware and IoT/OT

- Affiliates use IoT exploits.
  - Leaked internal documentation advocate the use IoT devices as entry points because those are often left unpatched, are difficult to manage and many organizations don't have full control over

them. They also encourage their affiliates to leverage botnets that already have dormant infections on IoT equipment.
    - CISA has released a report about Conti where they mentioned a tool that Conti uses called Router Scan to search for things like IP cameras and network-attached storage devices,
- Other examples of ransomware groups targeting IoT and OT.
    - DeadBolt, Qlocker, and Checkmate that almost exclusively target internet-exposed and local network attached storage (NAS) devices.
    - Ransomware groups are now targeting specific Mitel VoIP devices
    - Trickbot malware, which is frequently used by ransomware gangs to drop the malware on Mikrotik routers to use them as proxies for command and control communication.
    - Other examples: EKANS in 2020 targeted specific OT processes.
- There is a growing trend of leveraging IoT and IoT devices by ransomware to either gain initial access or gain persistence.

State-sponsored ransomware

- There is an increasing mix of ransomware activities with state sponsored activities.
    - Conti has publicly taken Russia's side in the ongoing war. They work on crippling the Costa Rican government,
    - There are Chinese threat actors that have been using several similar ransomware families.
    - There was an alert about North Korean actors that are using the Maui ransomware specifically against healthcare organizations.
- State-sponsored ransomware is different from RaaS large scale ransomware deployments
    - Don't have the infrastructure to distribute the decryption keys to victims
    - It doesn't have communication activators so it's either targeting specific organizations for specific state sponsored objectives, or it's a disguise for other types of activities,
- This is a growing trend that can lead potentially to large consequences
    - State sponsored actors have the funding and the means to cause greater disruption than just exfiltration or encryption of files. For instance, they could target operational technology devices and cause physical disruption.

State-sponsored malware trends

- Increasing wave of wipers being used for sabotage, as part of the ongoing cyber war, or for destruction of evidence, but basically being used as part of operations on victims.
    - Sometimes these are disguised as ransomware, but they are wipers, and this is done through obviously Incident Response Type motivations and so on.
    - They work typically by overwriting or sometimes encrypted files or overwriting the Master Boot Record or the Master File Table.
    - In some cases, and they're most often used on IT systems, typically Windows, Unix based, or Linux based distributions. Recently there was a specific wiper called AcidRain that was supposed to wipe SATCOM modems in Ukraine at the beginning of the war.
- OT specific malware continues to abuse insecure-by-design capabilities of OT equipment.
    - Industroyer2 leveraged OS-specific wipers and a dedicated module for the IEC 104 protocol for electrical substations and
    - INCONTROLLER toolkit has several modules for interacting with different types of ICS devices
    - Both found before causing too much damage.

- The third trend in state sponsored malware is the use of botnets that compromise 5g devices.
  - Cyclops Blink botnet developed by Sandworm group as a possible successor to VPNFilter. It was taken down soon after discovery, but researchers mentioned it has been active since at least 2019. It has been almost three years where that was operating without being without being discovered, and targeted specifically WatchGuard and ASUS routers

Hacktivism

- More than 100 groups have conducted cyber-attacks since the beginning of the Russian invasion of Ukraine.
  - Most of these attacks were distributed denial of service, either leveraging devices that are exposed online or directly the resources of some sympathizers of this groups that communicate with the group over Twitter and Telegram and are recruited to be part of these of these types of attacks.
  - There have also been data breaches, deployment of wipers, distribution of propaganda and so on.
  - There have been some attacks on critical infrastructure. Some electric vehicle chargers in the Moscow area we're not operational after the beginning of the war because they work remotely and were hacked by some activists.
  - Currently, there are more than 70 of those groups that are active. Most of them are either located in Russia or Ukraine. But there are also other countries such as western Cuba, Romania, Poland, and Portugal, and Italy that have active groups.
- Cyber wargames.
  - Killnet uses DDoS tools to take down websites of critical infrastructure companies in US/Europe.
    - ✓ The communication of these actions is typically over Twitter or Telegram.
    - ✓ Most of their attacks say that you should blame your own government for these types of attacks.
    - ✓ Targets all over the western world, including at least one airport in the US.

Ms. Moussouris asked a question on OT/ICS malware: Can you give a little bit more detail on "abuse insecure by design"?

Mr. dos Santos: The issue is that when a threat actor reaches the OT part of the network where those devices are accessible, typically the interaction with these devices doesn't require a specific exploit. Those devices have insecure-by-design functionality. They have protocols that are not authenticated or not encrypted. Oftentimes, they have hard coded credentials that allow this communication from threat actors to happen without specifically exploiting a vulnerability. They just need to know how to communicate with that device. We recently released an analysis of the industry and one interesting thing that we saw is that the implementation of the IEC 104 protocol used by industry seems to come from a publicly available GitHub repository where somebody can pick up just how the protocol works and plug that into a malware. And once you have access to the OT network, it's possible to cause physical disruption directly.

Ms. Moussouris asked: Have the manufacturers of the OT and ICS systems been notified about the design changes that they should probably try to make to make these attacks harder?

Mr. dos Santos: We had research that was published on the 21st of June. The main goal of that research was watching for insecure by design issues in OT equipment. We've worked with CISA to notify several vendors and ask them whether things will change or not. In our case, we have notified at least 10 major

vendors of things like that, not specifically the IEC 104. Protocol, but other protocols. There will be changes coming in at some point, but as it is now, OT is insecure by design and the main changes are around the network perimeter or the network boundaries to not allow attackers to reach those devices.

3. Conclusion

Mitigation

- Why is mitigation possible:
  - The majority of them are not fully automated.
    - ✓ Dwell time on the network changes from time to time, but it's typically at least a week. Ransomware infections, which have a short dwell time because the goal is to encrypt and then when encryption happens everybody knows that the attack has happened, but those typically lasts for at least a week between initial access and data encryption.
  - The second point is that cybercrime as a service, such as the ransomware as a service model, means that there are up to hundreds of very similar attacks happening simultaneously. Which means that if you're paying attention to the threat intelligence, you will know what types of organizations are being attacked and what can be done to prevent these types of attacks as well.
  - And the third point is that most of the tools and techniques that they use are well known. Except in the case of some state sponsored malware, that that has zero specific zero days or specific point images that are that are not known at the time of that an incident happens.
- Recommendations:
  - Extending visibility, monitoring, and patching, all the cyber hygiene, to IoT and network infrastructure, not just traditional laptops and servers and desktops, but moving to the whole attack surface of the network.
  - Plan for virtualization and storage servers and other types of devices as main targets as well, not just as entry points. Back up VMs and make up services infrastructure is extremely important.
  - Monitor activity organized for hacktivist behavior, like Telegram and Twitter, and just see what's being discussed.
  - Investing in DDoS protection against this type of hacktivist behavior
  - Network segmentation to isolate IT, IoT, and OT devices and limiting network connections to only specifically allowed devices.
  - In the case of insecure-by-design devices, use deep packet inspection, monitoring, which basically allows you on the network to understand what kinds of interactions are happening with those devices and potentially if they are malicious or not, where you're coming from, and what kinds of requests are being sent to those devices.

Takeaways

- The attack surface is increasing, it's not limited to residential laptops and desktops and servers. There is a whole world of IoT, OT, and network infrastructure, virtualization machines and so on that is being targeted.
- Cybercrime. hacktivists, state sponsored actors are all leveraging this increased attack surface and
- Mitigation should be based on threat intelligence and prioritize this increased attack surface.

The last slide is only about references, but unfortunately, you cannot read it. Hopefully I can share the presentation later with you but it's only some of the reports that we have produced and where this data is coming from.

Ms. Miller asked (via Mr. Scholl) about factors used to create risk ratings: Mr. Scholl clarified that on some earlier slides there were some charts that ranked items by risk. She is asking if it's possible for you to share what factors you use to create those risk ratings, and those sorts of things.

Mr. dos Santos replied: It's based on the business criticality of the device and what impact it can have if it's taken down. Which is based on the type of device. An IP camera has a specific criticality. A computer has another one, a cell phone or smartphone has a different one. So each type of device has a specific criticality that we assign to it, based on the vulnerabilities that we observe on those devices. The risk score is based on connections observed from those devices, either inbound or outbound to potentially malicious IP addresses.

Ms. Fanti asked about Insecure-by-design and public GitHub: I didn't quite understand earlier when you were talking about the insecure by design OT protocols, the public GitHub repository that you mentioned. Was that repository describing the protocol itself, or was it describing an attack on an existing protocol or none of the above?

Mr. dos Santos replied: The repository was an implementation of a tool that is a client for the specific protocol. There are several specific protocols for specific processes or uses, in this case IEC 104 which is specific for electrical substations and communication with other parts of the network. We had an implementation of this communication with this protocol which was taken and modified by the threat actors to carry out the attack. The issue with insecure by design is that once you are in a network and you can communicate with those devices using the IEC 104 protocol, there is no authentication required to do that communication. Any packets sent on that protocol, with a specific command, is processed by a device regardless of the specific authentication that has happened previously. That's what insecure by design means.

Ms. Fanti asked: Do you see a lot of open-source implementations being used to devise new attacks?

Mr. dos Santos replied: I don't have specific data to give you, but yes, that does happen. Often proofs of concept, details about a specific vulnerability, or even ransomware source code has been posted on GitHub and then repurposed. The Chinese threat actor who had at least five different ransomware families in less than a year, one or two of those came from repurposed source code that was leaked.

The Chair asked: What, if anything, is being done on the vendor side to fix the insecure-by-design issue and what can the customer do?

Mr. dos Santos replied: We always recommend network segmentation and network monitoring as the steps that can be taken when devices are insecure by design and cannot be patched. It's an unfortunate situation that we have in operation, technology and industrial control systems, protocols, and devices that are legacy. They were designed a long time ago are still implemented. The vendors are aware of the issue. They are trying move away with things they design, especially in this field where some devices are online for 20 or 30 years without being replaced or updated. There is a very big cost to replacing this item, but I imagine that we need to reach a point where those things are replaced.

Chair indicated that we have a quorum

MITRE ATT&CK Framework

Jamie Williams, Principal Adversary Emulation Engineer, MITRE

Introduction

Today I'm going to talk about the ATT&CK framework and give you a rundown on how we operationalize it. We strive for it to be practical. Something that the average consumer can use operationally or within their own security context.

Key takeaways

- First, ATT&CK enables us to see our threats rather than just looking at these names and categories. Whether it's county ransomware, or insider threats. One of the most important things you can do in this cybersecurity ecosystem is to start to decompose that threat. Know what you are defending against.
- Secondly, now that we can see a threat for what it really is the behaviors of threat actors, the actions they're taking, ATT&CK will provide an ecosystem to track, understand, and measure those behaviors.
- Third, knowing those behaviors, you can orient your defense recommendations toward those behaviors.
    - We can start to make threat informed decisions. Looking at why I do this mitigation. Why I need this product. Do I need a certain configuration? What does this provide against all those tough problems that we deal with day-to-day?

How do we measure Security?

- How do I know if I'm doing things right? How do I know if I actually need this product? How do I stack up against that new report? I'm seeing all this code on GitHub or an open-source capabilities. How would I defend that and where are my gaps or where are my vulnerabilities?
- This is something ATT&CK is structured to address at its core. Looking at those vulnerabilities, looking at those threats, and really decomposing them.

Threat-informed defense

- How can we start to tease out better metrics and a better perspective rather than just doing what we think feels right? Let's measure ourselves against our adversary. Let's look at how our adversaries, our threats, see these problems.
- Am I doing the right things in my prioritizations? Where do I invest my money to get the most impact against the threats that I am worried about?
- This idea of threat informed defense is going to cascade through, not only how ATT&CK is structured, but also how we apply it to problems.

How do we describe adversaries?

- There are different ways we can think about our adversaries and our threats. We can call those indicators of compromise or IoC. What the model is highlighting is not all meiosis are created equal. There's a scale. What really separates them is the cost or impact.

  - For example, hash values, shown at the bottom of the diagram, are easy to get. Anytime we've seen a threat report or a piece of malware, getting this hash value is trivial. If you are targeting and thinking about your adversaries at that level of abstraction, you're not imposing a lot of costs to them.
  - That's something that's going to vary operation to operation, host to host, victim to victim. in the model, as we work our way up this pyramid, we can impose more and more costs per adversary by thinking of them at different levels of abstraction. The peak which is TTP's.
- Think about the tactics, techniques, and procedures of our adversaries. I don't care what malware they use, how and what hashes were produced. Think about how they operate. What are their big objectives and their big SOPs? If we can start to undercut them and defend at that level, we're not only going to undercut them but really start to change the way they operate.
- Have these two ideas of threat-informed defense and TTP-based defense is going to be what cradles the rest of ATT&CK .

What is ATT&CK?

- ATT&CK is a globally accessible knowledge base of adversary tactics and techniques, developed by MITRE based on real world observations of adversaries' operations.
- ATT&CK helps us track and understand the behaviors we're seeing from our adversaries.
- Two things that make ATT&CK special:
  - ATT&CK is freely accessible. It has no paywall. There is no sign up. We publish it for free use.
  - It is based on real world observations. It's not MITRE or any other organization telling you some theoretical attack that may happen to you. It's looking at real activity and events that are publicly available and codifying it.
- ATT&CK also builds a common language as we're talking about different threats. We can use ATT&CK as a Rosetta stone, pointing everything back to this common language and this common model.
- ATT&CK has three parts:
  - Tactics: (The slide shows the "matrix" view of ATT&CK.) The top row is what we call tactics. These are based on the adversary goals or behaviors. Why are they doing what they're doing? Are they trying to gain and show access or trying to move laterally? Are they trying to impact the victim? Are they trying to steal credentials?
    - ✓ You might recognize a lot of these terms from a typical cyber kill chain. One of the things that separates ATT&CK tactics from that model is we're not prescribing a linear order. We are not saying adversaries start with the left and move right. We recognize that their campaigns can be all over the place and so can their behaviors. It can vary based on where they land and the bigger context in which it is happening.
  - Techniques: The other thing that separates a tactic from a typical kill chain are the columns coming down from those tactics, that we call techniques. For a given tactic there are various ways an adversary could achieve that tactical goal.
    - ✓ For something like initial access, they could try to use valid credentials to log in, they could use a trusted relationship via a service provider, or they could do a supply chain compromise.
  - Procedure. These are specific instances of what an adversary did to implement that technique.

Navigating Through the "Matrix"

- You can click on a cell and see descriptions of techniques
- Every technique has a common name and unique identifier that we can use to point back to this technique and apply it to different use cases.
- Also has a brief, easily understood write-up describing the technique to help the user understand:

- What is the behavior?
- What's the scope?
- What's being described?

Mr. Groman asked: Who the audience would be?

Mr. Williams replied: Any cyber practitioner who needs to understand this domain. The intention is to make sure that users don't need to have any background knowledge to use ATT&CK . We've seen high-school students to advanced professionals who can use ATT&CK .

- We created ATT&CK to support cyber operations. There is no prescription on how it should be used. It's more of a knowledge base of what adversaries are doing. This can be used to inform decision making. As mentioned earlier, we are curating publicly available intelligence and conceptualizing it into the model so it's easier to use.
- Example: User Execution: Malicious File:
  - In this context we're looking at procedures on user executions and behavior. Adversaries often rely on users to execute their malware. In this case, they may deliver a malicious executable or a malicious Word document with a macro and lure the user to click on it. That facilitates the technical goal of execution.
  - We're also looking at procedures of different adversaries and malware that have been associated with this technique. Under the "procedure examples," clicking on one of the attackers you get information on the malware used. All procedures are based on publicly available intelligence that can be found by clicking on the link. The link can include the names, the documents, and other indicators.
  - ATT&CK also provides information on defensive countermeasures.
    - ✓ Mitigations: At the bottom of every technique ATT&CK outlines defensive countermeasures in the form of mitigations to prevent this behavior from being successful or as successful.
    - ✓ We created high level maps to other control and risk frameworks with categories of mitigations like user training, execution, prevention, behavior prevention an end point, and other broad categories that we can use to prioritize and align our defenses against these behaviors.
    - ✓ Detections: These are data sources, data components, and descriptions telling us what data, as a defender, do I need to collect to see this in my environment, such as data about files, processes, or network connections. What within that data type do I need to worry about?
      - Creating a file?
      - Deleting a file?
      - Modifying a file?

Ms. Moussouris asked about collecting data and data types: Do you have any recommendations on how they should collect the recommended data and data types or pointers to tools for collecting them?

Mr. Williams replied: We reference things where applicable but with ATT&CK being so universal it's very tough. We recognize users are working in different environments. We try to find a certain level of abstraction that is more universal so everyone can see themselves in that.

Ms. Moussouris asked about built-in OS tools: Are the built in operating system tools you are looking at sufficient for collecting the data you are most interested in, or are they insufficient as packaged by the operating system manufacturers?

Mr. Williams replied: Varies case by case. In some cases, they are insufficient, but we also try to be mindful of scalability. There are a lot of ideas that could be possible but they're not pragmatic. In general

there's always room for improvement but we have a robust ATT&CK community. they contribute new ideas to ATT&CK and, we've noticed it circles back to vendors. There is a good feedback loop there.

Ms. Moussouris asked about application vendors: What about the layers of application vendors on top of the operating system; you can't boil the ocean of all apps, but is there a subset of most commonly deployed apps that you would recommend, to collect the data you need? How are you making the cut off?

Mr. Williams replied: That is something we try not to prescribe. In the spirit of ATT&CK, we try to enable users make that decision. For example, if I'm an OT provider and my environment is a lot of Linux and OT, in the spectrum of what TTP's matter to me is going to answer that question versus if I'm running an IT environment or mobile environment or hybrid environment. Threat modeling through ATT&CK is one of our big core use cases. There isn't a single path. You need to think about your situation and your organization, what threats matter to you, who would target you, and who has targeted you, and then mash that all together. The intent is to be threat informed and data-driven in your decisions.

ATT&CK Logistics

- ATT&CK is updated twice a year.
- the last release was in April 2022.
- There are 14 tactics, over 150 techniques, and over 350 sub techniques.
  - A sub technique is a broader category like supply chain compromise.
- Separate team focused on tracking CTI and mapping groups and software to this model including:
  - Open source tools, malware, or anything that can be publicly cited.
  - We have over 100 groups and 600 pieces of software that have been mapped into this model.

ATT&CK Spans Multiple Technology Domains

- There is the enterprise matrix:
  - What we typically think of as traditional IT systems,
    - ✓ windows, Linux, Mac OS, cloud, network, containers, infrastructure routers and switches, Virtualization
    - ✓ We have a couple tactics related to PRE,
- There is the mobile matrix that models the same TTP structure behaviors relative to handheld mobile devices, Android, and iOS.
- Our newest matrix is ATT&CK for ICS. It has the same structure but is for operational technology such as ICS SCADA; things that are not traditional IT systems.

Track and Understand Their Behaviors

- ATT&CK provides:
  - Context to understand TTPs. It is going to help us understand and contextualize behaviors. This slide showed an example a single command run by adversaries. ATT&CK helps us contextualize it right describing what it is doing. This gives us a reference model and a framework that we can use to tag behaviors that are relative to what this command is doing. This context can be used to triage these artifacts and to communicate to users. This slide shows an example of an alert from an end point security solution.

    The alert is using ATT&CK to communicate what telemetry it is seeing, and the technique associated with it. If I get an alert for a certain technique ID, I can go to ATT&CK and ask what considerations I should have such as: who's doing this AMA how do I defend against it, and all the rest of the rich content that's available.
  - A common language for describing and sharing TTPs. This slide shows an example of a joint security advisory about a schmoozing attack to share knowledge. It describes behaviors from a

county affiliate. It uses deuce, unstructured language to deliver this message but they are using a tactic to contextualize it. Where they describe a certain behavior, they can reference ATT&CK techniques, making sure that the message is being delivered on what they are trying to describe. They also used ATT&CK formatted tables showing tactics, techniques, and procedures of what is seen and what they are communicating to users.

Orient Your Defensive Recommendations Toward Behaviors

- ATT&CK provides:
  - A medium to measure yourself -
    - ✓ The medium is those threats we care about. The slide shows then example of a product built by the Center for threat informed defense. They mapped NIST 800-53 controls to ATT&CK. Anyone using that control can start to see themselves in ATT&CK across certain tactics and techniques. They can also dive into techniques of interest. This helps them see the bigger picture how to take those controls and think about the relevant behaviors. It helps them think about what those controls are covering, strengths, weaknesses, and what kind of factors need to be considered.
  - Means to identify gaps and prioritize operations
    - ✓ allows us to quantify things, see where we are, see where we need to be, and measure the difference
    - ✓ Operationally, we can start to see how to exaggerate our strengths, live in our weaknesses, and identify gaps that need to be filled
    - ✓ This slide shows a table from DHS as they outlined an incident response playbook. ATT&CK can be used to map high priority tactics and techniques to my environment,
      - where do they exist?
      - are there logs or events that need to be prioritized?

Conclusion

- ATT&CK is not going to be something that defines everything you need to do carry it it's not going to tell you how to solve cyber. It's not a silver bullet.
- ATT&CK will make you better at what you are already doing such as making day to day, month to month, and year to year decisions and understanding why you are making those decisions.
  - Looking at threats and thinking about what we are trying to defend.
  - Tracking and measuring how we are doing, what we need to do, and how we bridge that gap.
  - Making tangible, defensive recommendations and improvements.
- To reach us:
  - Email: attack@mitre.org
  - Personal email: jcwilliams@mitre.org

Discussion

Ms. Fitzgerald-McKay asked about the FiGHT framework: I know that there is an effort to develop an ATT&CK framework for 5G. Is there a lot of overlap between ATT&CK and FiGHT or do you see like FiGHT as an entirely new set?

Mr. Williams replied: Yes and no. This is not just FiGHT but other applications outside of the domains we've already modeled. There is a lot of overlap of behaviors that already exist in enterprise or mobile that are applicable there. But they also extend it with a couple of additional techniques. The challenge is that ATT&CK is built on what we've already seen and observed. 5G attacks, even if observed, are not widely reported. That challenge exists in other domains as well. As a result, it doesn't fit into the ATT&CK philosophy, because everything in ATT&CK is real.

The Chair brings up threat modeling: From a software development perspective, there is a technique called threat modeling which we use but which doesn't have anything to do with threats. Threat modeling is really a vulnerability analysis technique where you take a structured approach to try and look at what an adversary could do to system components. We talked about spoofing, tampering, repudiation, information disclosure, denial of service, automation privilege, etc. Is there any attempt to correlate the ATT&CK techniques that highlight and structure design approaches for threat modeling?

Mr. Williams replied: That's a big research priority for us and something we are actively investigating. The problem is anytime we do these multi-framework compilations, you end up producing a third framework with a different language and different ecosystem to deal with. What we try to do is borrow ideas from other frameworks, like STRIDE or the Diamond Model, that have great ideas and ingest that into ATT&CK and reference it back so we can keep this core principle. We have recognized threat modeling as the big domain where there's a lot of space for improvement, but we are being very cautious about producing a tangent or a deviation of ATT&CK or STRIDE.

The Chair recessed the meeting for a 12-minute break.

DHS Binding Operational Directive (BOD) 20-01 – Develop and Publish a Vulnerability Disclosure Policy, and the Known Exploitable Vulnerabilities List

Branko Bokan, Cybersecurity and Infrastructure Security Agency (CISA), Cybersecurity Division, DHS

Introduction

Mr. Bokan reminded the board about his talk at the last meeting on the GovCar Threat Modeling methodology and that they have successfully converted to MITRE ATT&CK.

CISA's mission is to protect our nation against cybersecurity threats. They are part of DHS.

The office Mr. Bokan works for focuses specifically on protecting the Federal Civilian Executive Enterprise.

CISA Binding Operational Directives

- In 2015 Congress gave us authority to issue cyber directives.
- There are two types of directives.
  - Binding Operational Directives
    - ✓ Issued from the secretary of Homeland Security to the head of an agency.
    - ✓ They are mandatory and federal agencies must comply with them.
    - ✓ Strategic in intent.
    - ✓ Usually require some type of ongoing effort on the agency side.
  - Emergency Directives.
    - ✓ Tactical in nature.
    - ✓ usually in response to major vulnerabilities, major attacks, major incidents, and major threats. Whenever we need agencies to act quickly and do something quickly.
    - ✓ The most recent ones were developed within approximately 4 to 6 hours of us detecting a major event.
    - ✓ usually give agencies 72 hours, sometimes longer, to remediate or take certain actions.
  - We have issued approximately 10 Emergency Directives since 2019.
    - ✓ The first one issued in 2019 was in response to a major, worldwide, DNS tampering effort.

    - ✓ The rest were in response to some type of vulnerability that was discovered and was being actively exploited and agencies had to patch immediately.
    - ✓ CISA does not have full visibility into agency enterprises. Even individual agencies don't have great visibility into their enterprise. Every time there is a vulnerability, there is a lot of work on the agency side to determine the prevalence of the vulnerability and all the different dependencies.
- BOD 22-01: reducing the significant risk of known exploited vulnerabilities.
  - Includes a provision for CISA to create and maintain a catalog of known exploited vulnerabilities.
  - When a new vulnerability is added to this catalog, agencies are immediately required to remediate this vulnerability in two weeks.
    - ✓ CISA has the option to shorten or lengthen that period, depending on the vulnerability.
  - For a vulnerability to be added to the catalog-
    - ✓ It must have a CVE label showing it is a recognized vulnerability.
    - ✓ We must have a high level of confidence that the vulnerability is being exploited in the wild.
    - ✓ There must be some type of remediation available.
  - The only acceptable alternative for remediation, is removal of affected assets from the enterprise.
    - ✓ If you can't remediate, if you can't update, we don't accept work arounds.
  - The intent is to be more proactive in vulnerability management and prevent the reactive response of issuing an emergency directive every time there is a new vulnerability.
  - Emergency directives are very expensive.

Vulnerability Management Programs

- Part of federal agency vulnerability management programs includes actively looking for vulnerabilities with certain Common Vulnerability Scoring System (CVSS) scores, usually higher than 7.
  - CVSS score > 9: Critical
  - CVSS score between 7 and 9: High
- There was no federal mandate that said you must remediate criticals or highs with one exception.
- There was a need to prioritize which vulnerabilities to remediate.
  - If you look at the vulnerability database since January 1999 there are close to 180,000 vulnerabilities with CVSS scores in the catalog.
    - ✓ 36,000 are categorized as high and almost 20,000 of them are categorized as critical
  - Less than 4% of all vulnerabilities ever get exploited.
  - Many of the vulnerabilities that do get exploited have a medium score or lower.
    - ✓ adversaries are chaining them, for example, they take 2, 3, or 4 vulnerabilities with CVSS scores of three or four, chain them up, and come up with one single vulnerability or a set of vulnerabilities with a CBS score that adds up to 15 or 20.
    - ✓ Heartbleed is an example of a vulnerability that was moderate but was one of the nightmare vulnerabilities to deal with.

CISA has already prioritized vulnerabilities

- The catalog of known exploited vulnerabilities helps agencies prioritize and mandates that they must remediate vulnerabilities that CISA determines to be known exploited vulnerabilities.
- Started with about 300 vulnerabilities in the catalog. Some were 10, 15, 20 years old.
- We originally said for all those old vulnerabilities with CVE numbers prior to 2021, you have six months to remediate. Anything new added to the catalog, you have two – three weeks.
  - The binding operational directive allows us to set that timeline depending on how we feel about that vulnerability.
- We don't only look at criticality. It depends on many other factors and what we know about that vulnerability.

- We started with the 180,000 vulnerabilities, 56,000 of them are categorized as critical and high. We used to have IGs on their tail forcing agencies to take actions and waste a ton of resources chasing after those critical and high vulnerabilities.
- Now, we have 787 known exploited vulnerabilities (KEV) that we added to the catalog. Of those there are approximately 20 most prevalent in the entire enterprise.
  - This brings the number down and helps agencies prioritize.
  - These are the ones we mandate you must re mediate first
  - you can use your risk assessment and existing risk management processes and procedures to determine what else needs to be remediated.

Reporting

- Every time we do a directive, we ask agencies to report back.
- We started with manual reporting this time because of that lack of visibility.
- We want agencies to report to their CDM tools and CDM federal dashboard so that we can fully automate this and have better cyber relevant tag visibility to the enterprise when it comes to known exploited vulnerabilities.
- We are asking agencies to do these reports as part of their FISMA quarterly reporting.
- The last report was on May 5th -
  - There were 654 vulnerabilities in the catalog.
  - 309 of these vulnerabilities were reported as existing in federal enterprises.
  - There were 1.5 million findings.
    - ✓ a host and end point can have multiple vulnerabilities, but each individual finding of a vulnerability is counted.
  - the top five CVE's accounted for more than 33% of all findings.
  - the top 20 were 60%.
  - The top 60 vulnerabilities accounted for 81% of all findings. This means if you remediate close to 60 vulnerabilities, you will reduce the number of total findings in the entire federal enterprise by 80%.
  - As of May 2022, the most prevalent vulnerability was in Chromium. This is not just Chrome browser; this includes Internet Explorer as well. This vulnerability accounted for a significant percentage. With very few updates, we could significantly reduce the exposure and the number of these vulnerabilities.
  - 65% of KEVs with a patch due date of March 21 and March 24, 2022, are in the Top 40 most dangerous software weaknesses.
  - 31% of KEVs are in the OWASP top 10: 2021.

Statistics

- Count of CVEs by Year Assigned (4/1/22)
  - Looking at all the known exploited vulnerabilities, most of them were discovered in the last three or four years. However, we still have vulnerabilities in the catalog that have a CVE label as old as 2002.
- KEV finding by CVE publish year and severity
  - Again, most of those vulnerabilities come from the last few years

Timelines

- Initially we gave agencies six months to start working with these vulnerabilities.
- As of today, agencies have three weeks for any new vulnerability.
- We are hoping that on October 1, 2022, a lot of agencies will start automatically recording their statuses through a CDM dashboard, so they don't have to do manual reporting.

BOD 19-02: vulnerability remediation requirements for Internet accessible systems

- Earlier I said there were no federal mandates requiring federal agencies to remediate criticals or highs, yet every single agency's IG was chasing them, auditing agencies every year and writing them up for not remediating their critical and high vulnerabilities.
- The only exception is BOD 19-02, issued by CISA in 2019.
  - Very narrowly focused on Internet exposed interfaces
  - we scan all federal Internet exposed interfaces through our cyber hygiene program and that binding operational directive does still require agencies to remediate criticals and highs within 15 or 30 calendar days.
  - there is no conflict
  - we are looking to merge the two directives as this one is three years old.

Discussion

Ms. Moussouris asked about staffing: There was a CVE that was withdrawn from the list. It was the wrong one. How have you been staffing and resourcing this program so that this does not happen again?

Mr. Branko replied: That CVE was not wrong. there was a CVE published in May by Microsoft. Unfortunately, the updates that addressed that CVE and another CVE that was not a critical known vulnerability broke the authentication.

Ms. Moussouris clarified: It was the wrong one. I looked at both. You withdrew the one that didn't have a breaking patch. It's CVE 26925 versus 26923.

Mr. Branko replied: He would look into that. This was a last-minute call we received from Microsoft.

Ms. Moussouris clarified her staffing question: There is no correction mechanism built in which is why you are now hearing about it in detail from me. That brings me to my question of how many dedicated staff do you have on this project to try and ensure that errors like that don't happen, because they are expensive errors. You are performing a centralized, authoritative prioritization. That is, in my opinion, quite dangerous to do. It's also something that other governments have experimented with and decided against doing because of the risk of miss-prioritization and accidentally giving the signal that this is more of a ceiling than a floor. It really does need to be a floor of compliance.

Mr. Branko replied: This is a limited staff. We are putting more stuff into it, and we carefully weigh all the options and risk of potentially breaking things versus having adversaries breaking into our systems.

Ms. Moussouris asked about effectiveness: Have you studied the effectiveness in reducing successful federal attacks as a result of federal enterprise compliance with patching the KEVs?

Mr. Branko replied: We have only had two cycles of reporting. We need to do a lot of improvements on data quality, but that is the next step. We are very interested to see the effectiveness.

Ms. Moussouris asked about planning for scale: The list is already nearly 700 vulnerabilities long and you haven't been doing it for very long. What is the plan for scale and perhaps retiring some items off the list? What's the plan for that in the future?

Mr. Branko replied: I mentioned that analysis and studies show that less than 3% of all CVEs ever get exploited. We are ready to go as far as we need to in terms of removing vulnerabilities. That is something we are still not doing, because once the vulnerability is exploited it is present. If you look at the federal enterprise, we are seeing see the ease going back 20 years that are still present in the federal enterprise. It

would be irresponsible to remove it if there are still systems that run software assets that are vulnerable to a vulnerability.

Ms. Moussouris asked about interagency communication and misalignment in timing that brings a whole other question as to why we still have Internet exposed federal assets?

My last question is, there is a misalignment between the timing that's required with this binding operational directive and some of the vulnerabilities on the list. You mentioned the newer vulnerabilities have a shorter time window to apply patches, whereas some of the vulnerabilities that are up to 10 years old, had been given six months to patch.

The first part of my question is why were the old, known, actively exploited vulnerabilities given the longest time to patch?

The second part of my question is that the Federal Trade Commission came out and noted that it was going to start taking more aggressive action especially for patching known vulnerabilities in the private sector. What was interesting was the example CVE they gave for one they would not tolerate being out there still was one that was given six months by CISA to apply at the federal level. I'm wondering what interagency collaboration is going on to make sure that FTC enforcement of the private sector does not exceed those recommended and enforced at the federal enterprise level?

Mr. Branko replied: These are great questions. The six-month provision was given only initially when we first published the directive. That was to accommodate legacy assets and legacy software. Based on our experience with emergency directives, in the past the number one reason federal agencies were not able to meet deadlines was directly related to legacy software. Software that is at the end of life and no longer supported by vendors, running mission critical systems. This was a hard decision and there was a lot of opposition, that's why we dropped it.

The binding operational directive in catalog is only mandatory for federal civilian executive branch agencies. However, the rest of the world is looking very closely at what CISA is doing, starting with the vendor community and the cybersecurity community. We can't really tell other agencies or other entities what to do and what to enforce.

The Chair asked about CVSS scores: My question comes from a different perspective. Is anything being done about working back from non-exploited vulnerabilities to see if there's a way to predict what the bad ones will be? There's no religious significance to CVSS scores. They are attempting to help organizations prioritize patching and, if they're not correlated with known exploitation it seems to me, they’re potentially out there just waiting. Are you trying to look at it from that direction as an alternative way to prioritize? I'm not sure if that's something NIST ought to be thinking about too.

Mr. Bokan replied: We've worked very closely with NIST, even within issuance of our directives, as part of our process. Matt is also part of that workflow. I'm not familiar with anything right now.

Mr. Scholl added: We talked a lot before the BOD went out. current CVSS data it does not have in its base score and exploitability metric. Organizations can tweak those scores based on the importance of the asset or exploit. To some extent, CISA is shaping the CVSS for the government as our internal vulnerability piece. we don't publish CVSS scores beyond the base scores.

The Chair added: I created the Microsoft vulnerability rating system, the one that goes back to 2002 and that was based on a combination of potential damage and exploitation as we perceived it across a general user world. We got feedback on it over time. It seems that the catalog of things that are being exploited

gives you a basis for looking back and identifying similar attributes and vulnerabilities that might be coming.

Mr. Scholl replied: We have not done that research. We have done research more on looking internally at trends overtime on CVE and CVSS, by product, company, and class. Looking at where clusters of prevalence of types of CVE exist. how see the existence changes over time through reporting. There are a lot of bad conclusions, and maybe one or two good ones, you could draw from that. We have not looked at it from the perspective of active exploits. A lot of that is internal to NIST, where the staff doesn't have visibility into that type of active threat activity like DHS does. That's when we partner with DHS.

The Chair clarified: What I'm saying is take the active exploits, see what they're telling you about the character of vulnerabilities, and then see if you can analyze the new ones coming over to predict which might be exploited.

Mr. Scholl agreed that it's a good idea.

Ms. Moussouris added: Microsoft created a new way of consuming the Microsoft bulletins, centered around the idea of predicting exploitability and that was called the exploitability index. It was added to the bulletin format in 2008. The idea was, it would give a score of the likelihood that reliable exploit code would appear in the wild within the next two weeks because, releasing a zero-day exploit is one way to drive exploitation. Another way is to release a patch because those patches will get reverse engineered, the vulnerability will be found, and the exploiters will go ahead and write their proof of concepts and exploits and malware accordingly. but we noticed our customers did not change anything they did. They didn't understand it. We noticed that they still prioritized based on the overall criticality scores. What are your thoughts about helping to predict which new vulnerabilities might have similar characteristics?

Mr. Branko replied: In terms of prediction, we are still struggling to enforce the 787 vulnerabilities and figuring out where they are. But those are things that we are working very closely with our partners, interagency partners, industry partners as well, including Microsoft. Those are all things we will need to think about.

Mr. Groman added: I don't think they're resourced for it. Someone above them must take this seriously and provide them with the resources they need, and the government needs to do this effectively across the government.

Mr. Branko added: Even with this relatively small number of vulnerabilities, it is still a lot of work for federal agencies to find and remediate them because there are a lot of dependencies out there, especially when you run legacy software.

The Chair commented: They ought to have a plethora of resources to patch them. Mr. Groman commented that isn't happening.

Ms. Moussouris mentioned: One of the things I cautioned CISA about before the binding operational directive for vulnerability disclosure programs took effect was this exact problem, that federal agencies are not staffed to keep up with vulnerabilities we already know about. What is the plan for helping with the over a million unmitigated federal endpoints? Is there some sort of plan to provide resources or recommendations? I know you don't have the power to give resources to these federal agencies, but how is that being addressed as this has to do with not just keeping up with the known vulnerabilities but any new vulnerability reports that are coming in as required for the VDP.

Mr. Groman commented: The question goes higher up, which is when is the government of the United States of America going to take cybersecurity seriously? we are going to need a lot more resources. Sites are given limited resources and it's much better than it was when I was in the White House but given limited resources, they're taking the steps they can. But it is shocking for many of us that the vulnerabilities are not patched. I think this is an important step. I think it's important for educating people at all levels of government. And I think what they're doing is critically important, but it is a tiny piece of a much larger effort that's required. it's binding, but not to the extent that someone can call and say, you have this number of vulnerabilities you have not patched so stop what you're doing and reprioritize cybersecurity to X, Y, or Z. I'm sorry constituents get angry but what you're doing is not secure. We don't do that.

Ms. Moussouris asked: What more resources could we, as a board, help recommend to support CISA in these missions and in considering some of these longer term questions?

Mr. Bokan replied: You're absolutely right, but this also requires a big cultural change. We have a lot of feedback objecting to us telling agencies how to do this. More importantly, agencies are saying this is great but our IGs are still auditing us based on this non-existing mandate that I have to remediate all my critical vulnerabilities and I'm being written up by my IGs so that is my priority. We are working with a big community and it's going to take time to change the culture, the way we think, but we are seeing a lot of progress and a lot of good things coming out of it.

Mr. Groman added: Your cultural point is incredibly important and people who have not spent time in various parts of the federal government may find it difficult to understand that culture is not going to be the same as at a major company. The culture is unique. A major culture shift is required, and it touches so many things. It can touch budgeting, and it can touch HR... that goes way beyond this mandate.

The Chair and Ms. Moussouris: Added that they hoped Mr. Bokan found the questions and feedback helpful. That the intent is to serve and if he thinks of anything else that the board should consider for recommendations, please let them know.

Mr. Bokan thanked them for all the questions and acknowledged that it's challenging.

The Chair recessed the meeting for a 54-minute lunch break.

FTC Patching Guidance

James A. Trilling, Federal Trade Commission (FTC)Division of Privacy and Identity Protection
Alex Gaynor, FTC Chief Technologist Office

Introduction

- James Trilling
  - Attorney with the Federal Trade commission's division of privacy and identity.
  - This is the agency that is most directly focused on enforcement, policy, and other issues surrounding data security.
  - I do a mix of investigations and enforcement-oriented work, and coordinate with our experts in our division of consumer business education that help create our business guidance, and coordinate with our Chief Technologist office.
- Alex Gaynor
  - With the FTCs Chief Technologist
  - Background as a software security engineer.

- I work with our case teams to bring a technologist perspective, but I work with our attorneys and investigators to look at factors, particularly in our data security cases, but also on other cases in our Bureau of Consumer Protection as well.

Data Security Enforcement

- We have a multi-pronged approach when it comes to data security issues, enforcement, policy work, and consumer and distance education.
- Enforcement-
  - section five of the FTC act prohibits deceptive or unfair acts or practices.
    - ✓ a deceptive act or practice is like a misrepresentation or a mission of material facts likely to mislead a consumer acting reasonably, for example, misrepresenting a security practice on a website, on product packaging, in a privacy policy, or elsewhere could be a deceptive practice in violation of section five of the FTC act.
    - ✓ Unfair act for practice is one that causes or is likely to cause substantial injury that consumers cannot reasonably avoid. That is not outweighed by benefits to consumers or competition.
      i. The important thing there for this group is that the failure to implement reasonable practices for securing consumer data is an unfair practice in violation of section five of the FTC act.
      ii. When we assess the reasonableness of data security practices, we are considering factors such as the sensitivity and volume of consumer information that a business holds, the size and complexity of the differences, operations, and the cost of available tools.

Improving Security and Reducing Vulnerabilities

- We look to see that a business has processes in place to identify and mitigate known and foreseeable vulnerabilities and risks.
- In addition to section five of the FTC act, we also enforce sector specific laws that include data security requirements. Those include:
  - The Children's Online Privacy Protection Act, which applies to some entities that collect personal information from children under the age of 13
  - The Gramm-Leach-Bliley Act, which we enforce against certain financial institutions that collect customer information
  - The Fair Credit Reporting Act which applies to consumer reporting agencies and entities that obtain information from them for purposes such as establishing a consumer's eligibility for credit employment, or insurance.
- Notably, the FTC is different from state Attorney General's or other state enforcement authorities in that we generally do not have a data breach notification law to enforce.
- We do enforce a health Breach Notification Rule against certain entities that are within the FTCs jurisdiction.
- Our policy work includes things like:
  - providing technical assistance and feedback to Congress or to the administration on proposed legislation or potential legislative recommendations,
  - providing comments to NIST, or other agencies on initiatives like the privacy framework, and
  - submitting reports to Congress and publishing reports and best practice recommendations on a variety of issues related to privacy and security.

Consumer Education

- With respect to consumer education, we actively engage in blog posts about security issues, or individual FTC data security and privacy cases.
- We also issue publications on topics like:
  - deleting information from hard drives before disposing of computers,

- recognizing, removing, and avoiding malware and
- securing home Wi Fi networks

Business Education

- Our efforts include things like:
  - issuing press releases about each FTC data security case,
  - public posting of the pleadings in the cases,
  - blog posts about the cases,
  - business blog posts related to security issues.
- We have a variety of business guidance publications related to security issues, such as:
  - Careful Connections: Keeping the Internet of Things Secure,
  - A data breach response guide,
  - A publication called Start with Security, which is a set of 10 lessons learned from the FTC's first 50 law enforcement security cases,
  - A series of blog posts called Stick with Security, provided more detailed guidance about the same lessons in the Start with Security publication,
  - We regularly collaborate with NIST and other federal partners on some of this business guidance, including for small businesses.
- Our consumer and business education publications are free to order and distribute.
- Our business guidance can be found at https://www.ftc.gov/business-guidance.
- Subscribe to a variety of FTC updates press releases, and blogs at https://www.ftc.gov/news-events/stay-connected.
- We have a tech blog post that provides guidance and messages, either from our technical experts or a combination of our technical experts and legal staff.
  - Some of the recent tech blog posts might be a particular interest to this group, including:
    - ✓ a January 2022 post warning consumers to remediate the log4J security vulnerability, https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2022/01/ftc-warns-companies-remediate-log4j-security-vulnerability,
      i. Intended to help raise awareness about the log4j vulnerability because of the ubiquitous use of log4j software in a wide range of systems found in consumer facing products as well as enterprise software and web applications.
      ii. That post provided links to
        - CISA guidance to help position stakeholders to check to see whether they use the log4j library
        - a link to help stakeholders update the log4j software package to the most current version
        - a link to additional guidance on mitigating the vulnerability.
      iii. Reminded readers that the FTC Act, the Gramm-Leach-Bliley Act, and other laws impose a duty to take reasonable steps to mitigate known software vulnerabilities and did state that the FTC intended to use its full legal authority to potentially pursue companies that failed to take reasonable steps to protect data from exposure as a result of log4J or vulnerabilities in the future.
      iv. The FTC goals in issuing that particular post included:
        - raising awareness of the seriousness of the log4j vulnerability,
        - helping to point people toward resources provided by CISA and by Apache and
        - gaining the attention of legal counsel and distance executives so that they would engage in dialogue with information security teams, engineers, and other relevant personnel regarding log4j, and the general topic of vulnerability management, patching and information security.

v. The FTC engaged with the team that worked on the Cyber Safety Review Board report about log4j, and we will continue to look for constructive ways to raise awareness of significant vulnerabilities in software and resources for addressing them then.

✓ May 22 blog post on the importance of effective breach disclosure, https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2022/05/security-beyond-prevention-importance-effective-breach-disclosures.

i. Recent tech blog posts about breach disclosures reiterated that good business incident response and breach disclosure practices are important parts of maintaining reasonable information security programs.

ii. That post pointed out that a breached entity that fails to disclose information to help parties mitigate foreseeable harm could violate section two of the FTC Act.

iii. In so doing that post discussed recent cases in which the FTC alleged that failure to timely notify consumers or other parties after breaches was part of what rendered particular businesses information security practices unreasonable and discussing cases in which the FTC alleged that breach entities violated section five of the FTC act by misleading consumers through public statements about the breaches.

- Explained that those types of facts can be the basis for a section five violation, regardless of whether any breach notification law require information to be disclosed.

The Chair asked: could you go into more detail about the patching guidance that you're referring to? There was some interest in the patching guidance and the information that you had provided around your patching guidance.

Mr. Trilling clarified: This a comment or a question about the log4j patching guidance and what else the FTC has provided in terms of guidance about patching. The rationale for highlighting the log4J vulnerability was because of its ubiquitous nature, the seriousness of the vulnerability, the perceived likelihood to exploitability, and because of the somewhat unusual nature of the vulnerability. We were certainly aware of other stakeholders' interest in making sure that businesses knew about that vulnerability and knew about it serious nature and had easy access to resources that CISA, Apache and others were making available to help entities know whether they were affected by the vulnerability and how to mitigate it. The general message in that blog post about the possibility of potentially facing an enforcement action for not having reasonable patching procedures in place and for not availing myself of information that is available about the serious nature of a vulnerability and how to mitigate it is consistent with past FTC cases and were part of the FTC validation that particular entities' activities only included the failure to have reasonable procedures in place to identify vulnerabilities and to apply available patches. That was certainly an issue in the FTC data security complaint against Equifax and has been an issue in other FTC data security cases as well. And it's been something that's been highlighted in business guidance pieces, including the Stick with Security publication that I mentioned.

Mr. Gaynor added: Another contributing factor and the motivation was the expectation that there was, or very shortly was going to be, widespread exploitation of the log4j vulnerability. I think in the cyber safety review report that came out this morning is this discussion on all the contemporaneous reporting at that time about the combination of ubiquity with the exploitability was a big factor.

Mr. Trilling added: We have frequent discussions with our colleagues at CISA, NIST, and other components of the federal government about how we can use our resources to highlight issues and to highlight resources that are made available regarding security. Past examples of that include blog posts highlighting some scanning tools that individual businesses can request from CISA. We've also put out education about the NIST cybersecurity framework, highlighting how the components of the framework match well against the types of lessons learned in that Start with Security publication, and in the types of

issues that the FTC has highlighted in individual data security cases, as practices that companies did not implement reasonably and therefore did not have reasonable data security.

Ms. Moussouris asked for clarification on reasonable vs unreasonable: Can you tell us a little bit more about how you gauge what is reasonable and what is unreasonable in making your decisions about these cases?

Mr. Trilling replied: We are looking for businesses to have processes in place, to be evaluating risks, and to be taking actions on an ongoing basis to mitigate known and foreseeable risks. Having practices in place to know what applications you're utilizing, what software you're utilizing, having a system inventory so that, when vulnerabilities are part of what's being discussed in the software community and in the larger community and are being flagged, you know whether the vulnerabilities are within one's infrastructure. Having processes in place to be able to receive updates and do evaluations of the feasibility of timely implementation of updates in a particular environment. The way that we weight these things as with any reasonableness inquiry is going to be fact specific, and I would go back to some of the specific criteria that I mentioned. What's going to be reasonable from organization to organization is going to vary somewhat depending on the nature of the data that could be put at risk, the size and complexity of our business and its data operations, as well as what information is out there to help identify vulnerabilities and what tools are out there to mitigate that. We're looking to have processes in place to be able to identify risks generally and vulnerabilities and then to take action to mitigate them in a timely manner.

Ms. Moussouris asked about timeline coordination between FTC and CISA: The example you gave of the log4j notice and the Equifax case with the $700 million settlement as an example of the kinds of enforcement actions. You mentioned that you coordinate with CISA quite a bit, but I asked in the last session why the FTC was out of step with CISA in the known exploited vulnerabilities list, where CISA was giving federal agencies all the way until May to patch Apache struts, but your warning to private industry of enforcement action came in January for that same vulnerability. Given that, we've seen a little uncoordinated communication there. What is the plan going forward for coordinating further with CISA, especially with regard to the extensive timelines they've given for some old vulnerabilities being patched in the federal space? It just seems like it would be hard to argue that the private sector deserves enforcement action when the federal enterprise is getting even more time to patch the same vulnerabilities.

Mr. Trilling replied: I don't think that blog post intended to, or did convey, a particular deadline for a particular organization? I think part of what you're asking, or part of what I'm inferring from the question you're asking is, would the FTC have acted against a particular entity that didn't remediate or patch log4j May and, if so, how would the FTC respond if the entity said Well, how could you have intended for us to apply a patch before May? There was guidance out there giving federal agencies until May.

Ms. Moussouris clarified: The guidance from CISA was instructing the federal agencies to patch Apache struts by May and you were using it as an example. And federal agencies were being given until May to patch a 10 year old vulnerability.

Mr. Trilling replied: I'm not sure what particular guidance you're referring to or what precipitated it, but the way we approach thinking about vulnerabilities, and if we're doing a reasonableness determination, there can be more of an issue than just patching. There may be other steps that an entity is taking or can take to reduce the likelihood of a particular vulnerability being exploited, even if a patch isn't available or if the entity isn't implementing an available patch for a given period. If, for some reason, the entity decides that it's not feasible to implement the patch, when we do a reasonableness evaluation, we would certainly consider steps short of a patch that an entity is undertaking to mitigate a vulnerability. What we wouldn't want to see is that an entity is doing nothing to monitor or review logs to isolate systems or

potentially remove applications or software that may have a vulnerability if that's a feasible step for the entities short of implementing a patch.

Ms. Moussouris clarified: my question was actually about coordination between CISA and FTC.

My last question is, how many resources do you have to pursue these types of investigations and what additional resources could we as a board potentially recommend helping support your mission?

Mr. Trilling replied: In the Division of privacy and identity protection, which is the group of attorneys and other professionals in the Bureau of Consumer Protection that directly focus on the investigations and enforcement, we have a team of about 40 which is primarily attorneys. We do have a technologist as part of that group. We have other technologists in our Chief Technologists office and in our low-tech group that also provides support on research projects on investigations and other related work. In the past year, we've been growing those ranks. The Chief Technologists office has grown significantly.

Mr. Gaynor added: We are at about 10 or 12 folks with a variety of different backgrounds. Unfortunately, I don't have the exact count. When asked in congressional hearings or in a congressional request, FTC leadership has explained how the FTC would and could use additional appropriations to grow technical staffing and technical expertise in particular. As I'm sure the board knows, many of the legislative proposals on privacy legislation or on security legislation include additional enforcement authorities for the FTC and appropriations for personnel.

Mr. Groman asked Ms. Moussouris to elaborate on her concern: The missions of FTC and CISA are very different, but if there was a disconnect in some way, what is the concern that you want to raise?

Ms. Moussouris replied: The guidance that was issued by FTC in January on Log4J was a very helpful warning. However, the example given is on the known exploited vulnerability list of CISA. That list is aimed at the federal government but shared with the public, so the public is using it as well and it is an important resource for the public. There is a disconnect between FTC saying, if you are in our jurisdiction as a private company, we're going to go after you for enforcement if you don't patch vulnerabilities like Apache Struts that led to the Equifax breach, whereas CISA is publishing these known exploitable vulnerabilities list that applies to federal government agencies, but it's giving federal government agencies even longer to apply patches for the same known vulnerabilities. The conflict I'm trying to highlight here is that the KEV is used by the public as a guide and the FTC is saying we're going to enforce at a stricter measure in the public sector than the federal government is holding itself.

Mr. Groman speculated: The FTC doesn't bring a case based on one action, one mistake, or one failure in security. In each case that was not patched but I would speculate that the entity had many other things that were wrong with the program in its totality, therefore, a law enforcement action was brought. But I would also say that requirements on the federal government or civilian agencies would not be dispositive for the Federal Trade Commission either so there certainly could be a conflict that we want to avoid. It is conceivable that the federal government might tell their agencies, you have a year, but in another context, for an entity in the private sector that has highly sensitive data, giving them a year would be entirely unreasonable. As a result they may not be in sync.

Mr. Trilling agreed: The blog post reference to Equifax did point to the failure to patch a known vulnerability as why it highlighted the Equifax case, but that's not the full extent of the FTC complaint about Equifax. In the FTC assessment of reasonableness or unreasonableness, under the law is always going to be fact specific. The application of what is reasonable is going to look different from one organization to another.

Supply Chain Risk Management – National Initiative for Improving Cybersecurity in Supply Chains

Jon Boyens, NIST

Revisions Process for SP 800-161 Rev. 1

- The foundational guidance was released around 2015.
- Drivers for the current update:
  - The Secure Technology act at the end of 2018
    - ✓ created the Federal Acquisition Security Council
    - ✓ told agencies to start implementing supply chain risk management
  - October 2018, we spent about six months going to different departments and agencies looking at who is doing what.
  - If agencies were doing supply chain risk management, we looked at what they were doing and what was working.
  - Some departments and agencies had been implementing aspects of supply chain risk management due to the Consolidated Appropriations Act.
    - ✓ Told NASA, NSF, Commerce, and Justice to look at moderate and high FIPS 199 information systems and anything coming from China.
  - Around 2020, we had a meeting with CISA where we listened and asked agencies what they needed to better implement C-SCRM.
    - ✓ Foundational to all the things that we updated in the publication
- Released May 2022.

Summary of Changes in SP 800-161 Rev. 1

- Our approach is to look at different key practices in terms of foundational, sustaining, and enhancing practices, starting with what every organization should be doing immediately.
- We integrated cyber supply chain risk management into broader ERM activities at the enterprise level.
  - Came from GAO report that was released in December 2020
    - ✓ Didn't find much activity
    - ✓ Wanted agencies to have a very tight understanding of their risk appetite and risk tolerance.
- Provides guidance on development of C-SCRM Program Management function.
  - We recognize a lot of agencies may have to do this ad hoc.
  - Typically as the tier one level of an organization in the OCIO, or finance, or somewhere where they have an ability to get information from different parts of the organization and fulfill some of those overarching ERM roles for supply chain risk management.
- Addresses Critical Success Factors
  - C-SCRM in Acquisition
    - ✓ This is relatively new. It was not in the 2015 version.
    - ✓ The Secure Technology Act, Federal Acquisition Security Council, departments, and agencies were asking for guidance on acquisition.
    - ✓ Worked closely with GSA.
  - SC Information Sharing
  - Training and Awareness
  - Key Practices
  - Measures and Measurement
- Includes templates to enable C-SCRM
  - Templates for:
    - ✓ strategy and implementation planning

- ✓ policy
- ✓ C-SCRM plan
- ✓ Risk assessment plan

  - Give an agency the ability to quick start their processes in managing supply chain risk.
- Removed a lot of information from the body of the document to try and streamline it.
  - Includes Appendix E on Federal Acquisition Supply Chain Security Act (FASCSA) of 2018 and Appendix F on EO 14028 Section 4(d), Software Supply Chain Security.
- Updated or new tables and graphics.
- Increases usability by creating audience profiles and user guides.

Appendices

- Dumped a lot into the appendices.
- A lot of our guidance is applicable to different stakeholders.
- We developed user profiles for different types of stakeholders and consumers so they can look at the specific parts of the document they wanted to look at.
- The four main appendices:
  - C-SCRM Controls
  - Risk Exposure Framework
  - C-SCRM Activities in the Risk Management Process
  - Templates
- Two additional appendices:
  - Appendix E: FASCSA
    - ✓ Specific to departments and agencies
    - ✓ provides guidance on the federal acquisition supply chain Security Act
      i. Established the federal acquisition Security Council
      ii. Supposed to be consuming a lot of the risk information from agencies, but when these departments and agencies are looking at doing risk assessments on products, services, or suppliers, we wanted a way for them to consistently have supply chain information they can hand back to the FASC.
      iii. Provides guidance on making the collection and that information aligned and harmonized.
  - Appendix F: EO 14028 Sections 4(c)/(d) Response, guidance for software supply chain security
    - ✓ Asks NIST to provide additional guidance on supply chain security
      i. by including this guidance in 800-161, we were able to get public comment on our guidance.
      ii. we had already been putting out a lot of information stemming from the executive order
      iii. we looked at all our different controls in the security measure controls and we pinpointed certain ones we thought should be flowing down to a tier one level of the contract base.
        - we felt that a lot of these security measures where applicable to all types of software, not just the critical software verification.

      iv. we identified different controls where software verification or testing could be utilized
      v. SSDF and attestations.
        - Our guidance took a low bar on attestation and first party self-attestation
        - For situations where an agency may need greater rigor in that attestation or conformity assessment, it points to 800-161 for our additional guidance.

      vi. We touched on some of the emerging concepts.
        - software bill of materials (SBOM).
          - We tie that in with the vulnerability management.
        - We also touch on open-source software and the enhanced vendor risk assessments

- targeting some of the gaps we saw in the SSDF where some of the practices don't really address foreign ownership, control, and influence.

Ms. Miller asks (via Matt Scholl) about FITARA: Efforts are under foot to rewrite FITARA. Have you been included in any of the FITARA rewrite efforts?

Mr. Boyens indicated that he has not.

- When we released 800-161 in May, we removed the content from appendix F and put it online with the rest of the executive order guidance that we have.
  - We did this because the public comments stated they wanted all our guidance in one place.
  - Since we're taking it out of the publication, it allows us to update it on a more frequent basis than we would typically with a special publication.

What's Next

- Errata update
  - There are a few changes, a few errors that we were unable to make because of our deadline.
  - There are some clarifications folks have been asking for on vulnerability disclosures.
- Quick start guide will get organizations moving.
- Move online
  - We have tools to make the consumption of our guidance easier to use and more consumable.
    - ✓ The Cybersecurity and Privacy Reference Tool (CPRT) allows users to navigate through the guidance more easily.
    - ✓ The Online Information Reference (OLIR).

National initiative for Improving Cybersecurity in Supply Chains (NIICS)

- In August 2021, the White House hosted a cybersecurity summit.
- Secretary Raimondo announced the NIICS
- In February 2022, we put out a request for information that included:
  - questions regarding this cybersecurity framework
  - many questions on this initiative
- We are still going through the feedback we received.
- In the next couple of months, we will probably be going out and having a more detailed conversation.

Ms. Moussouris asked about SBOM guidance: In your guidance around SBOMs, I'm wondering if you have some differentiation for producing SBOMs and consuming SBOMs? If you do, what does it look like?

Mr. Boyens replied: We recognize SBOMs have been around for a while on the developer side. We are focused on agencies consuming SBOMs, which is not an easy task when you consider the automation scalability, etc., we don't want to get out in front of the work that's going on out there. Lots of departments and agencies are saying this is still nascent. So that and, probably, the OSS will be the two main places we are going to be updating on an ongoing basis. Our guidance on SBOMs is paper thin because of all the work that's going on right now.

The Chair commented on supply chain risk management in the IoT world: When I think about supply chain risk management in the IoT world, I think about breaking the problem up into four components:

- Do you trust your suppliers?
  - All the secure development in the world it doesn't protect you from a malicious supplier
- Secure development

  - it can be fully cleared, absolutely trusted but if they're not building the software securely, you're still at risk.
- Product integrity
  - If you don't protect your source code, check the signatures, send out patches, protecting your repository, then the other things don't matter.
- Everybody whose software you're consuming has to do all the other things. It's recursive.

As I read 800-161 it seems those issues get conflated and some of them maybe get lost. From the perspective of someone who's trying to find out what am I supposed to be doing, or is my supplier doing the right things, the structure of the guidance is confusing. The executive order is excellent on product integrity and secure development, but again they sort of get conflated. This is an appeal for thinking about things in those terms and, maybe, structuring the guidance to help people understand that.

Mr. Boyens replied: I agree. We had a few meetings on how to structure, especially after the first draft, and that was one of the big comments we received. We received different recommendations on how to organize and it was largely dependent on who is providing those recommendations. It was difficult because we have different user groups using it within an organization. If we go online, we can have an easier time doing that and developing guidance that is specific to the different users and breaking it up. When we are talking about processes, there are different ways you can communicate that organization. You could do it organizationally, roles and responsibilities, or processes. But you can't do it all in one PDF. When we go to the quick start guide, I imagine there may be several. We would have something for the acquisition community, something for the OCIO shops, the operational side of it, etc., there are many ways of breaking it up.

The Chair replied: I agree that targeting things by audiences is a good thing and providing online access may help you do that. The breakup I just described is a way to fundamentally look at the problem you're solving. If you're not looking at the problem from a fundamental perspective, then there's risk of either including stuff that is unnecessary to anybody or putting things in the wrong place. I think it's important that you look at the problem you're tackling from the perspective of what is supply chain cybersecurity and what are the different pieces that compose it.

Mr. Boyens replied: I think when we go to those quick start guides, that will likely be a guiding factor.

Ms. Moussouris called for a clear delineation between vulnerability management and coordinated vulnerability disclosure: Steve mentioned a lot of concepts being conflated. One of the most common that I see conflated are vulnerability management and coordinated vulnerability disclosure. Including in your document very clear differentiation between these two practices would be good because they both affect supply chain, and they are related but they are distinct.

Mr. Boyens replied: We've been hearing that since we released it in May. When we do the errata, one of our goals is to clarify that. That community is active, and we don't want to get out in front of it.

National Security Memo on Promoting Untied States Leadership in Quantum Computing While Mitigating Risks to Vulnerable Cryptographic Systems (NSM-10)

Matthew Scholl, NIST
Natasha Cohen, CISA

Introduction

- Natasha Cohen

- I am the senior liaison to NSA
- Here as the representative of CISA as Chief of the Post Quantum Cryptography Initiative

National Foundation – NSM 10

- NSM 10 is the basis for the United States government strategy and actions on post quantum cryptography. Signed by the President in May 2022.
- Deals with post quantum cryptography and quantum technology across:
  - Research and development,
  - Supportive quantum technology, and
  - Quantum information science.
- Deals with both the federal civilian executive branch as well as national security systems.

Threat

- The basis for the threat is a crypto analytically relevant quantum computer that could break the modern asymmetric encryption algorithms that we rely on across our critical infrastructure, the US government, State, and local technologies.
- The timeline looks out close to a decade or so, but that still puts the information today at risk.
  - An adversary could collect information to be decrypted later.

NSM Timeline

- (Slide shows a timeline with milestones)
- CISA'S activities complement the rest of the interagency.
- We have been careful to coordinate with our interagency colleagues at missed, an essay, and OMB on the activities we are doing that fall into these four LOEs:
  - Risk analysis,
  - Focus, engagement, and partnerships planning,
  - Guidance and directives that we may push out in the next couple of years,
  - Engagement with our stakeholders.

NIST Crypto Standards – PQC Scope

- We are really talking about the asynchronous encryption technology.
  - Public key-based encryption.
- The rest of this is still good to some extent.
  - What the extent is, is still a question.
  - Symmetric algorithms are still good.
  - Hash functions are still fine. We have published guidance on quantum resistant hash-based signatures used in limited cases like code signing. Hash based signature can be a little brittle in its implementation. May constrain its use for some immediate use cases or places where it's appropriate to use. We will use our quantum resistance stuff in some of the other areas where more ubiquitous signatures might be needed.

NIST PQC Standards - Milestones and Timeline

- The important thing on this slide is the last bullet: the 2024 PQC standard publication.

What was Selected First

- **Mr. Scholl continued:** This is what I talked about yesterday.
- A primary key establishment mechanism and primary signature, and then some alternate signature schemes.

What is Next

- Standardization is our immediate next step
- NSM has a directive for us to work with CISA, other agencies, and industry putting together a project to help with the migration
  - That problem at the NCCoE is intended to help folks build out tools, techniques, practices, processes, and capabilities to:
    - ✓ Identify where they are using potentially vulnerable encryption.
    - ✓ Prioritize which should be transitioned first and help with expenditures.
  - Develop potential playbooks for risk-based acquisition and implementation of quantum resistant cryptography.
  - The quantum economic development consortium published "A Guide to Quantum Safe Organizations."
    - ✓ It outlines a lot, like quantum confusion, quantum key distribution, quantum randomness, and what quantum is and what it means.
    - ✓ It talks about the same things we talked about in the NSM around prioritizing and making decisions.

Request for help from the Board

- This requirement is something important and I'm going to need the boards help with:
  - "To mitigate this risk, the US must prioritize the timely and equitable transition of cryptographic systems to quantum resistant cryptography, with the goal of mitigating as much of the quantum risk as is feasible by 2035."
  - It's not quite "undo" everything. It's, find the important stuff. Make sure you know what the important stuff is.
  - The important thing is they put a date down, 2035.
- That seems really far away, but as we've been saying, there's a store and break issue, that's really the current threat. You have to think about cryptography as the last day you decrypt, not the first day you encrypt.
- According to research, it took 20 years to transition from DES to AES. That was a block cipher and not as prevalent as RSA. It was also starting in the 90s. There was a lot less tech out there.
- Estimate it took up to 20 years to transition from SHA-1 to SHA-2.
- We need to start these transitions and get people rolling.

When exactly?

- We don't know. Surveys of smart people say anywhere from 15 to 20 years as best guess
- The key thing is, we've transitioned from a theoretical issue to an engineering issue. Now, it's just about smart minds making it work.
- The vast majority of risk lies in the vendors. If we can get the companies that provide the foundation for IT security, IT management in critical infrastructure to transition, then updating technologies to supporting technologies, moving unsupported technologies out of our ecosystem, will get the vast majority of risk out of our ecosystem.

Deprecating Diffie Hellman and RSA

- One of the things NIST has not yet done, but we need to do, is decide on when to deprecate Diffie Hellman and RSA.
- As we go forward, once we've published his standards, we're going to need to work with you on making a transition date.
- **Question for the Board about AES block cipher key size.**
  - And AES 128 implementation is fine today, but after a cryptanalytically relevant quantum machine, it's about 80 bits left to you.

- We have not decided if we should recommend higher key sizes, and an AES key deprecation as well.
- So those are my two asks of the board, help me with deprecations from your perspectives, your expertise, your industry engagement, and your experiences.

Discussion

The Chair asked about industry partners: What are you hearing from industry partners about 2035?

Ms. Cohen replied: It's pretty varied. we have large financial companies, large IT vendors that I've been following this for quite some time. They're thinking about it, they're working on it, that's the engagement Matt was alluding to.

We have a lot of other industry that doesn't even know this exists. We really take that onto ourselves to put out as much as we can to raise awareness through NIST. we just put out an alert in collaboration with Matt's team to tell owners and operators what to do about this. We will continue to do that as they release guidance and plan a lot more industry engagement, both through our field forces and our alerts and the systems that we have there. That's the real challenge that we see.

I think the other place this is going to be a real challenge is in the OT space. Our technology is not easily updated. We have some plans and discussions on how we're going to engage that side of industry because that's going to be one of the sticking points.

The Chair added: Thinking about making the transition, the thing that scares me is critical infrastructure such as power or water. That seems like it's going to be hard.

I know the folks who know about it but they're the ones who are who are keeping SHA-1, except where there's some MD5 that still crept through.

Mr. Scholl replied: There's different degrees that we know about it. I think Natasha right that the finance sector is leading the way in this, and other sectors have different degrees of interest, anticipation, and preparedness. I think the big vendor community is on board and engaging with us as well. It's the OT space and our non-US partners that we need to communicate with.

The Chair added: It's also putting the effort in even with the big ones who are on board to sort of get all the last little bits out.

Mr. Scholl agreed: These things die difficultly, which is our DES lesson learned. to some extent, because this is a bilateral communication thing, you just stop allowing downgraded communication. That is a helpful potential forcing function we might have. I just made that up.

The Chair and Mr. Scholl both indicated a hope that the Social Security Agency will be one of the earlier transitions.

Ms. Fanti asked about transition time for AES: Are you expecting the transition time for AES, if you were to change the key size, to be shorter than the transition from DES?

Ms. Fitzgerald-McKay replied: In some ways it might be easier because it's the same algorithm, but the things that need to carry the key size are not prepared for that key size. So the transition isn't about the algorithm. It's about all the things that rely on it. So, I wouldn't imagine it's going to be terribly smoother.

Ms. Moussouris made an observation on forcing upgrades: I have an observation from experience in trying to force an upgrade of something that definitely needs to be upgraded but running into some realities of our economy being driven by consumer spending in large part could be affected by making things stop working as intended at a certain time of year.

My question for you is, have you considered making sure that the deadline for these transitions does not interfere with the biggest online shopping days? Because requiring a cut off where certain things won't work anymore during the major online shopping season could have a big economic impact.

Mr. Scholl replied: That's a good point.

Ms. Fitzgerald-McKay commented: I'm not convinced that if NIST deprecates Diffie Hellman on a specific date, things like Amazon will stop working.

Mr. Scholl replied: it implies that if a full transition has not happened, there might be sections of the economy that might not interoperate with other sections of the economy.

Ms. Miller commented (via chat): This will be a huge shift for the Defense Department as well.

Mr. Scholl stated: Essye, you are 190% Correct. Within the US government, one of the biggest stakeholders that we are working with is the Department of Defense, not just because of their huge IT footprint, but also because of their legacy and hard coded crypto systems that are not necessarily as agile. That was our SHA-1 transition lesson.

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

Final Board Reviews, Recommendations and Discussions

Steve Lipner, ISPAB Chair

The chair opened the floor to Board members for reactions, recommendations, possible follow-up letters, on anything that was heard over the last two days.

Lessons Learned

Topics for Future Meetings

- Societal impact of standards and the Customer Experience
  - Topic raised by Ms. Miller, The Chair, Mr. Groman, Ms. Fanti.
  - Thinking about the discussion around the customer experience regarding security of customer citizen access to government systems?
  - The societal impacts of technology decisions and how do we best account for those in our technical guidance?
  - Including the broader impact on society, our economy, our environment, and our country when many companies have all decided that bad cybersecurity is an acceptable risk for them.
  - The aggregated cost of privacy breaches: If 10, individually, small entities are all hacked, and all the data lands in Beijing or Moscow, any one of those 10 companies might think that my data is not all that significant. But when an adversary combines all that data, suddenly that has a very serious consequences for the United States national security and our economy.
  - What policy changes are needed?
  - Language changes?
  - How does an agency determine what is enough security or privacy protections?
    - ✓ Bug bar?
    - ✓ Risk tolerance?
  - Possible topic for October session.
- Risk Management - Accountability and Acceptance of Risk in the RMF
  - Topic that evolved during discussion between Mr. Groman and the Chair.
  - Who is making the decision to accept the risk?
  - Forcing mechanism.
- Resourcing
  - Topic proposed by multiple board members through discussion.

  - Resourcing regarding the exploited vulnerabilities and what people are doing.
  - Lessons learned to revise the severity rating system.
- Interaction and interoperability between CSF, RMF, and other frameworks.
  - Topic proposed by the Chair.
  - How they interact.

Letters

- Description
- No new letters were identified in this meeting.

Next Meeting

- The October 26-27, 2022, 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 4:15 p.m. ET.