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

December 4 and 5, 2019

American University Washington College of Law Tenley Campus
YUMA BLDG Room Y112; 4300 Nebraska Ave, NW, Washington, DC 20016

Board Members	**Board Secretariat and NIST Staff**
Steve Lipner, SAFECode, Chair, ISPAB	Jeff Brewer, NIST, DFO
Marc Groman, Privacy Consulting	Evie Petrella, Exeter Government Services, LLC
Brett Baker, Nuclear Regulatory Commission, OIG	Maggie McGary, Exeter Government Services, LLC
Jeffrey Greene, Symantec	
Brian Gattoni, DHS	

Wednesday, December 4, 2019

Welcome and Opening Remarks

Steve Lipner, Chair, ISPAB, Executive Director, SAFECode

The Chair welcomed everyone to the meeting at 9:06a.m., Eastern Time. He encouraged Board members to be interactive and ask questions particularly on subject matters people find difficult to express.

Welcome and ITL Update

Dr. Charles H. Romine, Director, Information Technology Laboratory, NIST

The Chair welcomed Dr. Charles H. Romine from the Information Technology Lab (ITL) at NIST (National Institute of Standards and Technology) to the meeting to update the Board on ITL.

There is a continuing resolution that funded the federal government at the beginning of this fiscal year through the end of November. Subsequent budget continuing resolution is funding NIST through December 20, 2019 while the House and Senate reconcile the appropriations bills they have passed to ultimately provide the funding budget to the President for his signature.

The Office of Management and Budget and the Office of Science and Technology Policy recently put out their priorities memo. NIST has strong support from both sides of the aisle for increased investment in Artificial Intelligence (AI) and in the area of quantum information science. ITL is heavily involved in both along with other partners across NIST. In both cases AI is a key driver for improvement in measurement science and quantum information science. The lead laboratory for these activities is the Physical Measurement Laboratory. ITL has been partnering with them for the last eight or ten years in mathematical research that supports quantum information science.

NIST is conducting fundamental research on the power of quantum systems. The ability to encode information in quantum states of matter, and to manipulate that information, has profound consequences. It is an enormous potential innovation that could change everything about information and about computer science and computing in general. Additionally, the introduction of even more advanced AI capabilities has the potential to have a similar impact on the economy and on the efficiencies and innovations across the spectrum.

NIST has been involved in election security for a long time in partnership with the Election Assistance Commission (EAC). Dr. Walt Copan, the Director of NIST, is also the chair of the Technical Guidelines Development Committee (TGDC), which is the advisory committee to the EAC for guidelines on secured, interoperable voting. The TGDC and NIST's work on the TGDC resulted in an update to the voluntary voting system's guidelines (VVSG). The Board chair inquired if those guidelines would be covered during the afternoon session which Dr. Romine confirmed.

The Privacy Framework was finalized after the latest comment period closed. The community broadly agreed with the Privacy Framework put out for comment. The comments were adjudicated and the framework has been submitted for interagency clearance. Work continues on risk management. NIST was a pioneer in orienting cybersecurity as a function to manage risk. We continue to promote that viewpoint which has been adopted universally.

We have two frameworks in cybersecurity; the original Risk Management Framework that was developed under Federal Information Security Management Act (FISMA) guidelines and the Cybersecurity Framework (CSF) for critical infrastructures that was developed under Executive Order and subsequently codified into statute. NIST recognizes there is still confusion about how the two frameworks relate to one another. Work is being done to provide better integration and awareness of the difference between the two frameworks.

The preliminary draft of the Privacy Framework was issued in early September. The public comment period for that ended near the end of October. The overall reaction was generally

positive. A final draft is out for review. There is a perception that privacy is a subset of cybersecurity and that managing privacy risk involves ensuring that you have good cybersecurity. Cybersecurity risk management is very effective at addressing some risks associated with privacy but, in fact, there are other risks that have to be managed in the privacy space that are not the result of any failure of cybersecurity.

The foundation of our work in cybersecurity started with the work we did and continue to do in cryptography. Enormous strides have been made in strengthening the cryptography team. There was an activity on post-quantum cryptography initiated a couple years ago. ITL received around five to six dozen algorithms submitted through a competition. We are now down to twenty-six algorithms remaining in two different categories. One is for the exchange of key information and the other is for digital signatures; nineteen and seven respectively. The cryptanalysis continues in consultation with the global community. Our expectation is that we will select more than one algorithm. We are in the middle of a similar activity in lightweight cryptography. Finally, we have undertaken a critical function of automating the cryptographic validation process. The community feels we need to be faster and to scale better. One way to do that is through automation.

We have a research program in ITL, Fundamental and Applied Research and Standards in AI Technologies (FARSAIT). Dr. Romine co-chairs the subcommittee of the National Science and Technology Council on Artificial Intelligence and Machine Learning. An activity was initiated between the sub-committee and MITRE, who manages the federally funded research and development center that supports the National Cybersecurity Center of Excellence (NCCoE). We collaborated with them on terminology and taxonomy for adversarial machine learning. They did an extensive literature survey and the draft is out for public comment for another few weeks. Our goal in the year is to build a test bed for AI at the center.

Draft Special Publication 800-207 on Zero Trust Architecture has been released. It is one of the most important documents to come out of NIST's cybersecurity program and certainly one of the most important documents addressing Zero Trust Architectures. The third draft of revision one of USGv6 which is the IPv6 profile for the U.S. Government was issued. Public comment ended last month. Work continues in the communications technology laboratory on the 5G cybersecurity. We continue to engage in the Internet Engineering Task Force (IETF) on core infrastructure issues and the robustness of the core infrastructure.

NIST has a key role to play in the evaluation of capabilities for algorithms in face recognition. There is broad participation across the industry. When testifying before Congress the biggest concern was with respect to bias. Congress was assured that NIST has been working on a comprehensive report on the demographic effects (age, race and sex).

The report is finished and should be released by December 20th, 2019. A Board member questioned that the biggest concern was bias, which Dr. Romine elaborated upon to say that it was not the only concern. Others are issues with law enforcement use of facial recognition, automatic enrollment of face recognition into systems, and the access to state driver's license databases by law enforcement which are out of NIST's scope of evaluation. A key driver for NIST is that they have the capability of answering how to quantify the level of demographic effects in those three dimensions.

The NCCoE has now developed 85 publications. There have been an enormous number of downloads for these applications. There is strong collaboration across the board with industry with 41 current partners. In the areas of IoT new projects are being started around healthcare and energy. We are hitting horizontally on work in ransomware. There is work getting ready to kick off in the area of patching. There are a lot of requests from the community to support the updates in the use of quantum resistant cryptography.

Dr. Romine briefly ran through a few examples of practice guides that have emerged from the NCCoE to provide a picture of the breadth of activity that's going on there such as: core issues like TLS server certificate management, practice guides on data integrity, trusted cloud infrastructure, property management system in hospitality sector, picture archive and communication in the healthcare sectors, asset management for the energy sector, to mitigating risks in small business and home IoT devices.

NIST will celebrate fifty years of cybersecurity research in 2022. There will likely be a major cybersecurity symposium.

Legislation Update on Draft Election Technology Research Act

Ms. Janie Thompson, Staff Director, Investigations and Oversight, U.S. House Committee on Science, Space & Technology

The Chair welcomed Janie Thompson to the meeting to brief the Board on the Election Technology Research Act.

The Science Committee began exploring election security after becoming aware of media coverage. Election security is an important issue across several committees within the agency, and the prospect of an insecure election was alarming.

Working collaboratively, the committee began planning a hearing to look at whether or not the media coverage was overblown or merited further investigation. They examined the types of attacks that took place in the 2016 and 2018 elections and found the types of election insecurities seen were largely attributable to technology rather than voter astroturfing influence operations on the internet. Some were classic cyber hygiene issues, such as spear phishing emails. Others were unique vulnerabilities associated with voting

systems like malware on an embedded computer hardware component on a voting machine. Also informing the hearing was a 2018 National Academies report called *Securing the Vote*. The report contains many actionable policy recommendations.

The committee began thinking about how to control the problem of election interference and whether there was a solution to move things forward. They approached the problem by separating the various stages of an election in sequence. There is the lead up to the vote and the pieces involved in that timeframe. When casting votes, there are also many variables; the same is true for tabulating the results of a vote.

The group immediately identified a gap in the Help America Vote Act. The Act has instructions and authorizations for NIST to research the election systems for the United States. It creates a definition of what the election system is and the definition is not all that inclusive. Ballot marking devises, optical scanners on new machines, and vote testing labs are included. NIST does not have a legal mandate under HAVA 2002 to test and serve registration portals, registration databases, local election websites, poll books, recording systems, ballot reconciliation methods, and maintenance and programing activities contracted by the vendors, which are important.

A hearing was held in June, at which Dr. Romine and four other individuals presented a strong case for more dedicated research and training in best practices in anticipation of the 2020 election. In November, Congresswoman Sherrill and Congressman Anthony Gonzales, both members of the Science Committee, introduced the Election Technology Research Act, ETR 4990 or ETRA.

ETRA has parallel features for NIST and the National Science Foundation (NSF). It expands HAVA to make sure NIST can look after the full range of technologies in election security when they do research and authorizes a clearer dedicated research program for NIST and NSF to do the work. It authorizes both NIST and NSF to create centers of excellence where they could give grants to universities to do some of the research. It instructs NIST to collaborate with the EAC to provide technical assistance by request to state and local election officials on how to implement best practices related to election cybersecurity. It authorizes GAO to issue a report on election security. None of this is compulsory; it is an authorization for NIST to conduct the research.

A Board member asked if there is a Senate counterpart to ETRA? Currently there is not, but they do their best to keep the Senate Committee posted on the progress.

A Board member noted that, based on what he's seen, the media is not blowing the issue out of proportion; in fact, the threat is extraordinary. To this end, they hope the message from the hearing is that the threat is as real as the press is reporting. Ms. Thompson agreed and stated that that message means more coming from any of the ISPAB members.

A Board member inquired if there are best practices that state officials are ignoring, or if the set of best practices are not uniform? It is both. There have been examples of a breach of some sort where everyone agreed on what happened; then there are individuals who willfully don't follow best practices in such events. There are also instances where the word doesn't get out, something that DHS is working on correcting, as are NIST and the EAC.

A Board member stated that merely introducing doubt about the voting system can potentially damage elections as much as actual intrusions. The Bill doesn't specifically touch on this but building confidence in the system builds a foundation to resist against that impact on voter confidence. Another member stated that the basic cyber hygiene issue being significant is a powerful observation so being sure to make this part of the communication is key. One best practice might be to stop connecting everything in the voting ecosystem. A member agreed with that suggestion, and another noted that usability and accessibility are important considerations.

A member asked whether the legislation will pass prior to the 2020 election? It's a possibility and a hope. The Committee will do whatever they can to work with the Senate and the House Administration Committee on whether it should remain a stand-alone bill or be broken out into individual initiatives—ideally bipartisan. Ideas under discussion are banning internet voting, mandating paper ballots, and risk limiting audits.

Brief on Election Security Threats and Vulnerabilities

Dr. Matthew Blaze, Georgetown University

The Chair welcomed Matthew Blaze to the meeting to brief the Board on Election Security Threats and Vulnerabilities. His work over the past 25 years has focused on computing system security and privacy. One of his main research focuses is on election security and election integrity.

Election threats and vulnerabilities are such broad topics that the topic is best approached by dividing into two sets of problems—voting systems and election management infrastructure. The first is the voting systems themselves. These systems are critical for elections and are the most visible representation of what an election is. Failures in voting systems were the main impetus behind the Help America Vote Act (HAVA) and have had repercussions, both positive and negative. When we think about the threat landscape for voting system components, overwhelmingly, the threat model has been retail level corruption, typically by candidates or their supporters who want to influence the outcome of an election.

Much focus is on voting machines used in precincts and direct supporting infrastructure including things like the design of ballots. Less attention is paid to other issues such as election management infrastructure. Election management infrastructure is far less visible

but equally important and vast in support of elections. The infrastructure is less standardized. It tends to be built and maintained, to some degree, by local election jurisdictions, which are typically counties and sometimes townships throughout the country. Voting systems are dependent upon this infrastructure for everything around the logistics of an election.

Election management infrastructure has a different threat, model tending to be less often the targets of a conventional retail election irregularity, and more often the target of disruption by hostile foreign actors. These disruptions are not as dramatic or easily detectible, and there is no clear path between individual software vulnerability and something like a voting machine and altering an election outcome. However, this infrastructure is equally critical for the integrity of the vote and the legitimacy of the ballot. The distinction of the two types of systems is Dr. Blaze' own way of thinking about them. The resource he relied on in forming his ideas is the *National Academy's Secure the Vote Report*.

Most of the national focus to date has been on voting systems. This is particularly true of one type of voting system called Direct Recording Electronic (DRE) voting machines, in which a voter interacts (typically) with a touch screen to make their ballot selections. Those ballot selections are stored in internal memory on the device, and then the internal memory is transferred to a counting center. There are a number of benefits to DRE systems which is why they were included in the HAVA. However, they have a singularly concerning architectural security property: If the software is flawed, whether through error or malice, we have no idea whether or not the recorded results are accurate.

Since the computer was invented, we have been struggling with the problem of software bugs and software misbehaving in unanticipated ways. Voting systems are no exception. Because of this, during the first 10 years following the passage of HAVA, the research community and the election integrity community focused on trying to find software defects and security vulnerabilities and fix them. Securing complex systems is no easier in voting than it is in any other system. In the case of elections, the remedies are much less available than they are in almost any other type of computing system, except perhaps for those that directly support human life.

A more recent approach has de-emphasized software and hardware integrity and emphasized architectural integrity of voting systems. It is intended to ask if the design of the voting system is one that can tolerate inevitable defects and allow recovery from software defects that are likely to be present, particularly in the highly stressed, malicious environments where elections take place. The answer is that there are some architectures that you cannot do this with, particularly DRE voting machines, but there are other architectures you can do this with.

This work has largely been inspired by a short paper by Ron Rivest, later revised by Rivest and Wack, called *Software Independence in Voting Systems*. It says we should design systems to be independent of the software they run on, such that an undetected software failure in the voting system can't result in an undetected error in the outcome of the election. The most prominent existing system that allows for this independence is the optical scan paper ballot. The voter votes on a piece of paper reminiscent of the SATs, fills in a bubble sheet with their ballot selections, then that ballot is fed into an optical scan reader that's used to create the tallied votes, but the system retains that piece of paper, which is a reliable artifact of the voter's intention.

A second game-changing result has been presented by Phillip Stark at Berkley, a statistician who described an efficient post-election audit technique called Risk Limiting Audits, in which a statistically rigorously defined sample of the recovered ballots are compared by hand to the computer interpretation of those ballots as part of the tally. This is a practical way of achieving software independence that has the side effect of relying on technology that exists, and in fact existed before HAVA. Unfortunately, the systems are not yet used in every part of the country. There are still jurisdictions that use DREs that are incompatible with risk limiting audits or reliable risk limiting audits. Risk limiting audits remain the exception rather than the rule.

The fact remains that while we know how to achieve software independence in voting systems, it is not a solved problem in practice in U.S. elections. It's an achievable goal with essentially off-the-shelf technology that exists. A lot of attention given to voting systems can be summarized by asking the question why are we not doing this step? Part of the answer is a political issue. The funding to purchase new equipment to conduct risk limiting audits is largely absent. Jurisdictions lack the training and experience to implement it.

There are questions on standardizations of the physical security aspects of these systems, particularly the chain of custody of ballots, as software independence depends on reliable artifacts. There are other questions such as disaster recovery. Questions of voting systems that can be rolled out in the face of disaster, displaced populations and so on, are largely an open question that we don't have good standards and infrastructure for. Finally, there's a public confidence question. Voting systems have been well recognized to be unreliable for the past 20 years. If people don't trust their voting systems, they may not vote.

Election management infrastructure systems are largely ad hoc. There are roughly 5,000 – 10,000 voting jurisdictions in the United States. In most states, they're counties, but in some places, these are towns or townships. Election management infrastructure includes what are called election management systems, which are generally provided by the vendor: the back-end software used to support particular voting machines, as well as for provisioning the voting machines and tallying results. There are other components to

election management infrastructure, including the management of voter registration databases, poll book data that's sent to precincts, data that is used to check voters in at their local polling place, systems that are used to report the tallied results, and other systems used to communicate with the public. These are all systems maintained by counties.

Failure of any one of these components has serious implications for elections. Many election management systems are running on desktop computers. Those computers may not be dedicated to particular election management tasks. These may be ordinary workstation computers that are running email, web browsers, and word processors, and all of the other vectors where software compromise exists. All of these functions have properties a threat actor needs. Many jurisdictions are outsourcing some functions to contractors or to the vendor of their voting system. They are generally communicating heavily by email, and they are generally allowing some sort of remote access by outside contractors and vendors to these systems. We know with a strong degree of certainty that in the 2016 elections, voter registration databases and other back-end functions in the individual counties were targeted.

A member asked about the obstacles to some of these solutions. Dr. Blaze replied that there are a number of obstacles—especially the fact that local election offices' resources are often sparse. The National Academy's Report estimates that one-third of election jurisdictions are small and have no full-time election staff. In every case, the local government budget for election administration is competing with other local government functions.

A member pointed out that they think market failure is a term that is often associated with the voting marketplace. It's unfortunate that there are proprietary systems with no interoperability standards for such machines. If you need to replace a voting machine, you're stuck with replacing everything all at once or going back to the original vendor. Another member added that there is a huge legacy hold on voting systems.

Brief on Microsoft Safeguarding the Election Program

Ms. Ginny Badanes, Director of Strategic Projects, the Defending Democracy Program, Microsoft
Mr. Dave Leichtman, Cybersecurity & Democracy Strategist, Microsoft

The Chair welcomed Ginny Badanes and Dave Leichtmann from Microsoft. Ginny Badanes runs a strategic projects team within the Defending Democracy Program at Microsoft and Dave Leichtman works on Ginny's team. They've both been working with Microsoft's Defending Democracy Program for about five years, and thanked the Board for inviting them to share a bit about a technology and source code called ElectionGuard.

Microsoft's original election-technology-focused efforts were not exclusively focused on security. They were working with campaigns on how they could leverage new and interesting technology that pushed the edges a bit. Microsoft felt an obligation to help protect future attacks on our democracy after the 2016 elections.

The Defending Democracy program is a non-partisan, global effort with three main pillars: election integrity, campaign security, and disinformation defense. The team is comprised of about 10 people. The impact of the team is actually much stronger because they have such wonderful cooperation from colleagues across academia, government, and industry. The organization that this initiative falls under is customer security and trust, so the focus is more on providing value to not just Microsoft customers, but essentially to global democracy with respect to protecting elections against both private and public actors.

Dr. Josh Benaloh, a senior cryptographer at Microsoft, has been focusing on verifiable secret-ballot elections since the mid-1980s; his 1987 doctoral dissertation focused on this exact topic. When the Defending Democracy program started, the team worked closely with Dr. Benaloh who advocated the idea of end-to-end verifiability and to make it commercially viable and available. The election companies who provide software and services to election authorities don't have research and development

budgets. The market wasn't necessarily aware of the technology or how it could help voters so they weren't asking for it either.

ElectionGuard is the result of the program's efforts to date centering around Dr. Benaloh's research. From a technical standpoint it offers end-to-end ballot verifiability. The non-technological goal of the initiative is increasing voter confidence. There are a few different paths for how ElectionGuard can work. The current version is out on GitHub. New code will be released soon to add additional administration features, but the basis of the code and the cryptography behind it is available. Microsoft's expectation and hope are that vendors will take the code and put it into their new and, in some cases, existing systems. ElectionGuard is not something Microsoft will sell. Their hope is that the open-source components will be adapted by various vendors and researchers to eventually get the technology to the voters.

ElectionGuard enables government entities, news organizations, human rights organizations, or anyone else to build additional verifiers that independently certify that election results are accurately counted and not altered. The resources available on GitHub today include a working verifier as well as the specifications necessary to build an independent verifier.

One way to verify ballots is to "spoil" (encrypt) them then provide a mechanism for voters to "unspoil" (decrypt) their ballots after verifying the information is correct, prior to

casting their vote. There are various options for when the voter could decrypt and verify before casting. Another aspect of ElectionGuard is audits around privacy. There's a question when performing a risk audit on whether or not to make the details of the entire record public. If you make them public, in theory, there could be some privacy and coercion risks involved. If you don't make the full cast vote record public, then the audit is happening behind closed doors. ElectionGuard can take the cast vote records, encrypt them using the ElectionGuard encryption, then make the entire tally public. One can see the vote details without revealing the details that could compromise privacy.

A Board member asked whether, in this scenario, the audit is done against the ID that's issued to each individual at the poll rather than the full vote. The voter never had a tracking number issued or identified, so they wouldn't be aware of that ID number nor would it be traceable back to an individual—the audit would be on the cast vote, not the full cast vote record. Mr. Leichtman offered another way to describe the process by comparing the voting process to a bank deposit process. In a bank safe, each safe deposit box has two keys. The bank has a key and the account owner has a key. While the bank officer isn't able to open a safe deposit box without the owner's key, the account owner still knows the additive sum of what's in the box. With regard to an election scenario, there would be a process that the election officials and the agreed upon trustees have to go through, but it will amount to taking a secure USB type device and plugging into a computer. The Microsoft team is working on administrative tools that will make that whole process simple to record, and something for which they don't have to create additional work for new systems.

A member asked what they mean when they refer to "trial election"—something like professional society elections or trials with real government elections? Trial elections will be small municipal elections. The advantage of trying with small local elections is that many use paper ballots so there are paper ballots to match against electronically cast votes.

A member asked whether Microsoft has had requests for testing with optical scanners. They haven't had the chance to trial with optical scanners but it would be a great test scenario, as voters can clearly match the unspoiled paper ballot to the votes on the screen and could indicate whether to accept and cast or to spoil for future verification. Microsoft is partnering with Columbia World Projects to do tests run by pilots to measure voter confidence – does it move the needle at all; how do voters respond.

There has been some hesitancy from some vendors whether ElectionGuard will work, whether anyone will want to buy the technology, and whether voters will find it valuable. Microsoft thinks that's a fair question, which is why they want to do some pilots and additional testing with real voter data. A member asked whether anyone overseas is using ElectionGuard. Nobody has used it but there has been a lot of recent interest expressed.

Government Accountability Office (GAO) Update on Report 20-256-T VA and Other Federal Agencies Need to Address Significant Challenges

Mr. Greg Wilshusen, Director, Information Security Issues, Government Accountability Office (GAO)

The chair welcomed Greg Wilshusen to update the Board on Report 20-356-T VA and Other Federal Agencies need to Address Significant Changes. The report was presented as testimony before House Veteran's Affairs (HVA) Subcommittee on Technology Modernization. The report highlights the status of information security costs for the Federal Government. The audit focused on the 23 agencies covered by the Chief Financial Officer's Act within the Executive Branch.

Civilian Federal agencies spent approximately 6.5 billion dollars on IT security related activities in FY 2018. That represented about fourteen percent of their total expenditures on information technology across the 23 civilian agencies. The analysis did not include the Department of Defense which spent over 8 billion dollars on cybersecurity – which is more than the 23 other agencies combined. For the hearing, they mentioned that the VA spent 3.8 million dollars, which is about 8% of its IT expenditures. A Board member inquired if the 8% has been consistent for the VA over time? Mr. Wilshusen noted he did not know but for the 23 civilian agencies the percentage is increasing. A Board member inquired if the government is getting better at security for the increased spending? Mr. WIlshusen replied that Federal agencies are still very much challenged in protecting their systems and information. Overall, federal agencies are still very much at risk. A Board member inquired how the numbers compare with the private sector? GAO has not performed that analysis.

Federal Agencies continue to report large numbers of security incidents, although the VA has reported fewer incidents in recent years. Over the last three years between 30,000 and 35,000 incidents have been reported each year by federal agencies to the Department of Homeland Security (DHS). In FY 2018, the vast majority of the estimated 31,000 incidents reported to US-CERT were due to improper usage, which speaks to how users are not complying with their agencies' policies. 27% of incidents were categorized as 'other' which are instances where the agency was unable to identify the threat vector through which the incident occurred.

OMB has identified any incident classified as 'other' as a high priority because it indicates a lack of agency awareness and ability to investigate and catalogue incidents. There are a couple of initiatives underway to help minimize and improve agencies' capabilities through the Security Operations Center Maturation Initiative, and also with the exfiltration detection programs. A member asked if GAO did a sampling of what is in the 'other' category? In 2014 a report was prepared looking at cybersecurity incidents and concluded that 68% of incidents reported to US-CERT were not adequately documented to determine

the incident that occurred. Agencies' capabilities to detect and investigate these incidents are improving, but they are still limited.

GAO will be issuing their report in February 2020. They have changed their reporting mechanism and process. In addition to speaking to the five core security functions of the Cybersecurity Framework, GAO always issues two reports. One is publicly accessible and one is not. The one that is not public contains detailed technical vulnerabilities or technical recommendations that are only made to the agency. Copies are provided to the agency, congressional requestors or recipients of the report, and other entities with a viable and vested interest. The agency copy, or 'Limited Official Use' (LOU) copy might contain upwards of 160 recommendations. They now issue the LOU report first and delay issuance of the public report for several months. This change allows for an agency to take corrective actions and resolve vulnerabilities. It incentivizes the agency to take corrective actions immediately. This change has elevated the attention level of the top leadership.

A review of the findings section of the annual FISMA evaluation reports issued by the Inspectors General (IG) was provided. FISMA requires the IGs to evaluate effectiveness of their agencies' information security programs using the IG FISMA reporting metrics to facilitate the process. The metrics include about 66 questions, which they call metrics, that are aligned with the Cybersecurity Framework or security functions. IGs for 18 of 23 CFO Act agencies determined that their agency's IS program was not effectively implemented during FY 2018. The majority of these 18 agencies are at level three or below for the majority of the core security functions.

A member asked if maturity levels are given to the core functions or to each policy and procedure? A rating is assigned to the questions in the FISMA report which informs a rating to the core security functions. The IGs then use those core security functions to inform the rating for the overall information security program. A member asked if there is any correlation between an agency at level four or five and having fewer of the proportional share of the estimated 31,000 incidents? There could be some correlation as to the number of incidents relative to the effectiveness but it's not absolute. It would be interesting to see if there is a correlation between being at level four or five and reducing incidents proportional to size, complexity, etc.

A Board member inquired if they are evaluating the privacy sections as well? Not as part of the FISMA reports that have been issued. It is something to look for going forward because there is a privacy element to it. They do look at privacy concerns during other evaluations and engagements. They have looked at privacy aspects related to privacy impact assessments; whether the agencies issued records notices for systems that they were developing, and making sure they were assessing those systems for type of information, and whether there are privacy concerns associated with that information. Another aspect

looked at was privacy breaches.

Most CFO Act agencies had significant Information Security (IS) control deficiencies over financial reporting in FY 2018. Six agencies with material weaknesses in IS controls include: DOD, DHS, HUD, OPM, USDA, and VA. Twelve agencies had significant deficiencies and six were without significant deficiencies include – DOE, DOI, DOJ, NRC, NSF, and USAID. Material weaknesses are the more severe kind of weaknesses that could impact the reliability of information presented on the financial statements. Significant deficiencies are those that are still significant enough to be reported to those charged with governance for the organization.

The administration developed key milestones and targets for modernization across agency priority goals which includes, as far as the security initiative, the key milestones and targets monitored through the CFO FISMA report section. Most civilian CFO Act Agencies, including VA, have reported meeting many cyber targets. All of the agencies are supposed to meet the key milestone metrics by 2020.

Briefing on Election Equipment Security Requirements

Ms. Gema Howell, NIST

The Chair welcomed Gema Howell to brief the Board on Election Security Requirements. Ms. Howell is the NIST lead for the cybersecurity efforts to develop standards for Voluntary Voting System Guidelines (VVSG) and is the co-chair for the cybersecurity and public working group.

The 2016 general election attacks included data exfiltration from voter registration systems, phishing election officials and voting system vendors, doxing of political campaigns, and attacks on backend non-tabulation systems. The threat model is expanding to nation-state phishing attacks on supporting election systems, and misinformation. NIST was mandated in the Help America Vote Act of 2002 to provide election technical support to the U.S. Election Assistance Commission (EAC). NIST develops the standards, requirements, and guidelines for the VVSG. They look into interoperability, provide research and assessment for human factors, and test methodologies. Additionally, NIST is charged with accrediting the test laboratories that perform the testing to the VVSG. NIST also provides additional best practices and often points to the Cybersecurity Framework.

Under the EAC there are three advisory groups: standards board, board advisors, and the Technical Guidelines Development Committee (TGDC). The TGDC is chaired by the NIST Director, Walter Copan. Ms. Howell, along with her team, develops the draft requirements for the VVSG and shares them with the public working groups to gain feedback and thoughts on requirements. Once the requirements are completed they are provided to the TGDC for review and approval and will then be recommended to the EAC. The EAC shares

them with the two additional advisory boards. Finally, the requirements will go out for a public comment period and then the EAC makes the decision to adopt.

The VVSG is used to federally certify voting systems. The certification for each state is not mandatory. The hope is that states test to NISTs requirements. A Board member inquired how many states actually require VVSG certified voting systems. Ms. Howell did not have that number. The majority of states are aware of the VVSG and somehow incorporate them. The VVSG requirements do not cover election management infrastructure such as voter registration, candidate filing, and campaign databases. The Direct Record Electronics (DREs), optical scans, and ballot marking devices are in scope of the VVSG requirements.

The current structure of the VVSG has been broken down into three parts: 1) Principles and Guidelines, 2) Requirements and 3) Test Methods. The intention of the principles and guidelines is to be shared with election officials in plain language so they can speak to what their voting system is being tested to while understanding the different capabilities and goals. There are fifteen principles and guidelines. The requirements get into the low-level guidance for manufactures and test laboratories. The test methods and test strategies specify how the laboratories test to the requirements. They recognize there are overlaps among requirements. Some of the security requirements might affect the usability and accessibility sections.

In 2007 a set of recommended VVSG requirements was put out. The NIST team reviewed those requirements as well as the current VVSG 1.1 requirements, conducted a gap analysis, and updated the requirements. They added a few items as well. They also reviewed what happened since the last iteration of the VVSG. The first version of VVSG was in 2005 and version 1.1 was put out in 2015. The team wanted to make sure they included some of the security innovations that have happened since that time. Additionally, they took note of security innovations in the voting system space. Software independence was taken into consideration as well as risk-limiting audits and end-to-end verifiable systems.

A member asked if something were to happen on Election Day, where does the process fit into planning and system development? Ms. Howell reported that it is not covered because the VVSG is focused on the system itself. Process and things that happen outside of the voting system are out of scope. However, the cybersecurity framework profile they are developing covers the whole election infrastructure. While working on the cybersecurity framework profile, they conducted a workshop that election officials, vendors, and test labs attended. They discussed high priority areas and concerns about accomplishing their missions.

A Board member noted that the expectation is that there will not be a computer scientist or computer engineer at every polling place on election day. If something does go wrong who

needs to be available the day before, the day of, or day after and at what level of expertise? Ms. Howell noted such concerns are covered under the profile work. One of the mission objectives is getting the right staff at the polling places. There was a long discussion around the amount of training that needs to happen because a security expert will likely not be available. It is important to at least make the staff aware with the information available and secured up-front.

Ms. Howell turned to discuss the development of an election infrastructure profile under the CSF. A workshop was held in August with people from the election community to discuss their mission objectives and the high priority goals they want to accomplish. Once everyone identified their mission objectives and goals they prioritized them under each function in a category of the cybersecurity framework. NIST worked with people in election IT to then prioritize objectives and goals those down to the CSF subcategories. They now have a draft and have developed a baseline profile to cover the entire election infrastructure where they have highlighted the high priority security expectations. The next steps are to point to the latest informative references so everyone has the necessary information. The profile can be used to support a self-assessment or as an example for others to create their own profile and analyze what their expectations are for their election infrastructure. The NIST team has been working with DHS's election infrastructure subsector, which consists of the government coordinating council and the sector coordinating council. A Board member suggested doing a pilot or a demonstration under NCCoE to get down to a level of specific examples.

Public Comments

No requests for public comment were received.

Meeting Recessed

The meeting recessed at 4:04 p.m., Eastern Time.

Thursday, December 5, 2019

The Chair opened the meeting at 9:07 a.m., Eastern Time.

Brief on the NIST Privacy Framework

Ms. Naomi Lefkovitz, NIST

The Chair welcomed Naomi Lefkovitz to brief the Board on the NIST Privacy Framework. The goal is for the framework to better communicate both inside and among organizations on privacy and managing privacy risk. A preliminary draft was released in September followed by a second comment period. Currently, they are in the final stages of work on version 1.0. The goal is to release version 1.0 by early 2020.

A Board member inquired if Ms. Lefkovitz found any significant changes between the draft for comment and the final draft? No, the comments did not reflect any significant changes. Most of the comments reflected a desire for more guidance. More clarification around the alignment between the Privacy Framework and the security framework was requested. People's positions varied in that some wanted more overlap versus others wanting no overlap.

Ms. Lefkovitz noted that they are trying to clarify the different sources of privacy and cybersecurity risk so that one can apply and implement better solutions. They have tried to level the privacy framework up to recognize cybersecurity incidents rising from the loss of confidentiality, integrity, and availability. On the privacy side they look at the occurrence or potential occurrence of problematic data actions. Where the risks arise from data processing, that's being conducted to achieve mission or business purposes that overlap, is really cybersecurity-related privacy events. It's those kinds of problematic data actions that are arising from some kind of loss of confidentiality, integrity, and availability. That is what they are trying to show within the overlap of the two risks.

Privacy risk is usually about focusing on problems that people (individual, group level, or societal level) experience. A problem arises from data processing that an individual can experience as a direct impact. They experience different types of impact which can manifest in customer abandonment, noncompliance cost, harm to reputation, etc. The Privacy Framework helps bring privacy into alignment with other risks that organizations are managing. By making the problems that individuals can experience more visible, and helping organizations understand them through typical risks that they manage, the hope is to strengthen privacy programs and protections and ultimately allocate more resources for privacy through the enterprise risk management process.

An example for clarification was requested by a Board member. One example are smart meters. People were rejecting the smart meters because the information that they were collecting was so granular. A Board member provided the example of the President's Precision Medicine Initiative which would have created a database of biometrics, DNA, medical records, sociodemographic information, etc. It became apparent that the FBI wanted the database for all kinds of reasons which sparked numerous concerns. These are the issues around uses of information being transparent, open, honest, ethical, and responsible, but they are completely different from securing it. The Board member continued with a thought on AI collecting data elements that go into machine learning. Inferences on that data creating new inferences onward which then produce an outcome. That is not a cybersecurity issue but rather a privacy issue.

The privacy risk assessment goes to the heart of the Privacy Framework. How do we optimize the beneficial data uses and minimize adverse consequences for people? How do

we manage the risks? Are you managing adverse consequences? Are you helping people to understand the new uses? That is why NIST came out with privacy engineering objectives that are complimentary to confidentiality, integrity, and availability which they have tried to reflect through the privacy framework document. It's about managing the predictability. The assessment helps to address the beneficial uses or whether the risks outweigh the benefits. By going through a risk assessment process, an organization can demonstrate the reasoning to get to a decision.

Appendix D was created during the discussion of the draft to dig further into some of the key privacy risk management practices. Risk assessments are one piece but it's also important to identify key stakeholders as well as assign roles in risk management. Create a diverse workforce when it comes to managing, understanding, and identifying privacy risk. The privacy engineering objectives help organizations think through capabilities they need in their system, define privacy requirements, select controls, and implement and assess them. Monitoring is included in Appendix D as well.

A Board member commented on having commonality or greater harmonization across different things at a granular level: how do you classify data, tag it, label it? That falls to better informing decisions on how you are going to protect it. A workshop on data classification was recently held. It's looking at the foundation of how we understand, classify, and tag data in a way that would help make more informed security and privacy decisions. Ms. Lefkovitz participated in the workshop and tried to make the important point that a risk assessment is an ongoing process. We don't want to classify something and then forget it. What is sensitive in one organization might not be for another.

A Board member noted that problems will arise with things like public disclosure of government and administrative data. We have seen two data sets published that on their own didn't seem sensitive. Analysts put the two together and were able to identify incredibly sensitive things about government employees. That is a bad decision about throwing data out to the public without understanding the consequence of putting them together.

Turning to the value proposition of the assessment tool, NIST believes that this framework supports organizations in building customer trust by supporting ethical decision making and product and service design for deployment. It optimizes the beneficial uses of data while minimizing adverse consequences for individuals' privacy. The framework can support organizations in fulfilling their current clients' obligations as well as helping them to future-proof.

The Privacy Framework structure, based on stakeholder input, was aligned to the Cybersecurity Framework. The core provides an increasingly granular set of activities and

outcomes that enable an organizational dialogue about managing privacy risk. The core functions include: Identify, Govern, Control, Communicate, and Protect. Profiles are a selection of specific functions, categories and subcategories that the organization has prioritized to help manage privacy risk. The implementation tiers help an organization communicate about whether it has sufficient processes and resources in place to manage privacy risk and achieve its target profile.

The thinking was to meet each organization where they are today. Overlaying the functions of the two frameworks over the Venn diagram shows what the principal functions are that help manage the different aspects of privacy and cybersecurity risk. There is a lot of flexibility and different ways to use the frameworks together. The conversation, both internally and with stakeholders, will continue on how to better align the two.

A Board member noted that privacy and security have to collaborate. If they do not collaborate both the privacy and security functions suffer as does the trustworthiness of any system or program. A lot of actions within security are data tagging which is integral to privacy. Communication between the two areas is critical and should be reflected in the document. Ms. Lefkovitz confirmed that encouragement of collaboration is reflected in the document itself. The process of developing profiles is where they see the collaboration happening and encourage it. There are hypothetical use cases to demonstrate how the collaborations work. A Board member suggested the publication of use case examples in collaboration with organizations who use the framework.

Ms. Lefkovitz noted that one of the next steps for the Privacy Framework is to build a roadmap. A Board member would like to see data tagging in the process. Ms. Lefkovitz confirmed that it is built into the technical standards. Another suggestion was to include a piece about inventory. It is included; however, they could emphasize it more. Ms. Lefkovitz noted that they are now shifting to adoption of the framework. There are organizations such as Equifax that are currently giving the preliminary draft a trial run. Staffers on the Hill have been briefed and sent drafts throughout the year. The main focus now is about getting the framework implemented and then maintaining it.

Cybersecurity Solarium Brief

Ms. Val Cofield, Senior Director and Lead for Task Force 3

The Chair welcomed Ms. Cofield to brief the Board on the Cybersecurity Solarium. The Cyberspace Solarium Commission stood up through FY19 based on the National Defense Authorization Act. It is a hybrid commission composed of fourteen commissioners with four sitting Congressional branch commissioners, four sitting executive branch commissioners, and six private sector commissioners.

The commission was tasked with looking at a holistic cybersecurity strategy that would

help protect the nation in the event of a significant cyber-attack. To conduct the work, the commission was split into four groups. The first task force, Task Force One, looked at the DOD strategy of defend forward and resistant engagement. Task Force Two looked at the resilience of the U.S. cybersecurity infrastructure. They looked at how to better bolster and improve cybersecurity resilience domestically. How do we better partner with the private sector, how to share information, and what are the impediments to sharing right now? Task Force Three, led by Ms. Cofield, looked at cyber norms and non-military instruments of power. The fourth group was called a directorate. They looked at issues like emerging technologies such as how will 5G and AI affect the work the other three task forces were doing? The task force reports are being consolidated into one final report.

An event was held last October where each task force issued a report on their research and came up with some high-level recommendations. The fourteen commissioners were present as well as a group of cybersecurity experts. They were presented with two scenarios. One scenario was called a short, sharp event where a series of catastrophic attacks happened. The attendees were asked to utilize their recommendations in light of creating a response to such an event, helping respond to such an event, or preventing the catastrophic event from occurring. The second scenario was called a slow burn. A slow burn is considered to involve constant issues such as those that are going on in cyberspace right now. The attendees were again requested to explain how their recommendations respond to either help prevent it or help negate that event. After the meeting the commission synthesized the good discussions.

The commission is now unifying across the task forces, taking the expert input, and coming up with a final report and list of recommendations. The goal of a lot of their recommendations will be to put them into law. The deadline for finalizing the report is around March 2020 which would hopefully be in line with when the FY2021 NDAA would be completed. Because the commission is actively negotiating the recommendations there are not a lot of specifics Ms. Cofield can discuss. As soon as they are able to share specifics, they will. Congressional hearings regarding the report will be in April.

A Board member asked Ms. Cofield to speak further regarding the scenarios. The slow burn scenario, which the country is currently facing, is where they have not been able to have an impact in deterring a significant cyber event. Trillions of dollars in IP theft and interference in elections are the kinds of real-life scenarios that are similar to the situations they presented. The short sharp events were scenarios in which adversaries are using third parties to help create significant events that occurred domestically.

There is a sense in the commission that the government needs to work better with the private sector. The private sector owns 85-90 percent of this infrastructure. How do we better partner with the private sector to defend ourselves? They are looking at cyber

insurance as well as intelligence sharing. Additionally, the commission is working through a wide spectrum of recommendations.

A Board member inquired if any of the scenarios related to how our society is totally dependent on commerce over military connectivity. How would we operate, short-term or long-term, if everything stopped working because of a malicious actor? Those scenarios are under active discussion. They are looking at whether regulatory reform with certain critical infrastructure sectors is needed. Additionally, they are looking at the U.S. government – both from the legislative and executive branch perspectives. They hope to have recommendations for both branches.

A member inquired about the scope of private sector participation or interaction. It has been extensive and especially with the critical infrastructure sector. There is extensive agreement that the government needs to work better with the private sector.

Update on the NIST Cybersecurity Program

Mr. Matthew Scholl, NIST
Mr. Kevin Stine, NIST

The Chair welcomed Matthew Scholl and Kevin Stine of NIST to update the Board on the NIST Cybersecurity Program.

NIST is embarking on an update to the NICE Cybersecurity Workforce Framework, NIST Special Publication 800-181. They will continue to offer expanded use of the framework, not just in the federal space but in the contract space as well. It is voluntary in other areas of the community such as private industry and beyond. The National Initiative for Cybersecurity Education (NICE) Strategic Plan update will be initiated as well.

NIST had a big year in the cybersecurity and privacy risk management space. The enrichment of risk management also includes the workforce. NIST recently announced the Advancing Cybersecurity Risk Management Conference to take place May 27th and 28th 2020 at NIST. The focus of the conference will be on cybersecurity and privacy risk management.

NIST has launched their online informative reference program, with respect to the NIST Privacy Framework, which started under the umbrella of the Cybersecurity Framework. The references extend into the privacy space and beyond. The references provide more science and rigor to the mapping process based on subject matter expertise from authoritative owners of other resources that are available. A workshop was held early this week at the NCCoE to lay out the mapping methodology. There was significant vendor participation. The value proposition for the community is that NIST can build this repository and continue to grow a set of resources enriched by an authoritative subject

matter expert that has validated the mapping. It brings greater standardization and consistency.

A member noted that flawed mapping is an enormous problem. What are the next steps? The methodology is out there but they want to continue to get organizations to exercise the methodology in parallel with producing mappings and then share with the broader community. Simultaneously, NIST will ensure they are mapping their own resources so that folks can see how the cybersecurity and privacy frameworks have been structured and that the risk management framework resources relate. It will enable them and the community to help with greater automation, so the mappings can be linked to specific controls. You can begin to build that relationship from the highest order of rule and regulation.

NIST is currently in the process of selecting a new director for the NCCoE. NCCoE continues to see very good engagement from industry partners. Zero Trust will continue to be a priority area for NIST. They have had a tremendous relationship with the Federal CIO Council and other parts of the United States Government (USG) and industry. A workshop was held to rally around zero trust. It was an opportunity for NIST to receive feedback on their next practical and actionable steps.

Data security is a priority theme at the NCCoE. There are ongoing projects in data confidentiality broken across five functions of the Privacy Framework and a parallel structure focused on data integrity. A workshop was held on data classification last month with much of the discussion focused on acknowledging that this is a very hard problem. 5G is another area where NIST has been involved with standards and specifications. There has been significant engagement with industry on what areas would be most appropriate and efficient for NCCoE to take on with industry help in order to advance the 5G experience with respect to cybersecurity. They expect within the next month to have a mission track description and some ideas for what they can do from a 5G perspective.

The goal at NIST is to make cryptographic transitions, which will occur over the next three to five years, as transparent as possible to end users. NIST continues to see significant changes to the standards and protocols, as well as the mathematics in use, on many of the encryption technologies. NIST is leading many of these changes both nationally and internationally. The timeline that NIST and the global community are on for developing and deploying their quantum resistant encryption standards is still relevant. They are looking at 2022 to 2024 for having those standards complete. Round two selection finished cutting from 69 to 27 candidate quantum resistant algorithms. They are now moving into round three. NIST is looking at the cryptanalysis and performance of each algorithm to determine if they are strong and agile. They are working with international partners to do hybrid beta testing implementations to garner feedback. With the announcement of Google's claim that they have quantum supremacy and have implemented an effective quantum circuit using

some of their applications, people are starting to get energized and want to buy something quantum safe. NIST is telling them to wait until the standards work is complete. Once the work is finished something will be codified that commercial industry can start to build from. It will be interoperable and strong. The gap space in an availability of a standard has shown that NIST needs to put out some guidance on assisting folks with transition such as how to plan, manage the risk, and be ready for a quantum transition.

The plan is to have something specific enough to allow organizations to make the right decisions at the right times in implementations. Quantum machines will factor a key at a time. There will be a period of transition that will be acceptable, depending on what and where the encryption is protecting. Part of the guidance will be to find out where vulnerable encryption is and what it's protecting and from that prioritize. Protocols will be different and easier to identify.

Last month NIST held a workshop with stakeholders on lightweight encryption and the needs for an algorithm that's not as heavy as something like Advanced Encryption Standard (AES). They learned that a single algorithm will likely not work for small gate size, small CPU, and small power. The stakeholders asked that they go back and look at the potential optimizations for each of the different constrained uses.

Updating and modernization continues on the crypto-module validation program (CMVP). The program has shifted to as much automation and black box testing as possible to speed the roundtrip time it takes for them to provide trust assurance in a cryptographic module. The standard was changed from a defense only to an ISO standard which will allow wider use on an international scale for commercial markets. This will open up NIST and U.S. algorithms for wider use across an international scale.

An Artificial Intelligence (AI) program has been initiated across ITL looking at different aspects of artifical intelligence. NIST is very interested in security and privacy in AI, particularly the use of AI in security and security of AI. There are a number of research projects going on such as looking at AI applications for things like vulnerability discovery, software bugs, and the ability of AI to identify patterns and anti-patterns in software. A member asked what makes AI different from other IT? In IT its clear; there's an exploit or misconfiguration that can allow for a hostile-intended actor to compromise. In an AI algorithm, there are more things around how it could be misused. These are potential integrity issues that have to be mapped within a data set, as well as the context and the engine.

NIST maintains the National Vulnerability Database (NVD). The database is very much privacy focused on the vulnerabilities, but it is NIST's ability to standardize, express, and find vulnerabilities in the NVD using the common vulnerability scoring sets which is

applicable for how NIST would capture vulnerability in AI use. NIST will look at updating some of our core standards and machine expression of vulnerabilities to capture things like, 'what is a vulnerability in artificial intelligence machine work?'. They are also looking at machine work capabilities and testing and performance where they have large data sets of information behind them that could be used for training as well.

Turning to identity management, NIST will have a conference in a couple weeks on access control in multi-cloud environments. They will look at whether there are capabilities and technologies that allow for smooth yet secure access control mechanisms that could cross multi-cloud environments. The identity team will be looking at that research technology.

Time is both a critical and almost invisible utility from which many of our infrastructures and trust extensions extend so NIST is very interested in the security and reliability of time sources. This is a research project going forward. ITL will be working collaboratively with DHS, NOAA, NASA, and the emerging areas of the commercial space infrastructure as well.

NIST held a DC Crypto Day bringing together many local universities who are engaged in mathematics and encryption technology as well as local industries and other stakeholders. They reviewed where they were in encryption technologies, plans for the future and then collaborated around different challenges folks are having in research as well as applying encryption. An International Crypto-Module Conference will be held in April in Bethesda, MD.

Brief on Zero Trust Networks

Mr. Alper Kerman, NIST

The Chair welcomed Alper Kerman to brief the Board on Zero Trust Networks. Mr. Kerman is the Project Manager for the NIST Zero Trust Architecture (ZTA) Project. The American Technology Council released the IT Modernizations Report around May 2017. Following, the Federal Chief Information Officers Council (CIO) chartered a steering group on Zero Trust and software defined networking. The Federal CIO and NIST hosted a workshop in October 2018 after which a project to do research on ZTA was initiated. The effort is driven by a Federal initiative. NIST is a leading partner in the technical work involved. NIST's goal is to provide general guidance on ZTAs, or trust architectures, for adoption in the federal government as well as perform gap analysis on existing technologies in the industry. All of the work was compiled into one document and Zero Trust Architecture: Draft NIST Special Publication (SP) 800-207 was put out for comment in September 2019.

Draft NIST SP 800-207 is a product of a collaborative effort between government agencies and was reviewed by the Federal CIO Council. NIST wanted to introduce a language to give a vendor agnostic conceptual view of what Zero Trust is and what ZTA should be. All of this is described in the document at an abstract level. They discuss deployment models for

ZTAs, use-case scenarios, and guiding steps for migration toward ZTAs. The public comment period closed in November 2019. They are currently incorporating the feedback into the document. They will introduce a few more sections to NIST SP 800-207 and it will likely go back out for a second comment period.

The document was the project deliverable. One of the bottom-up approaches on the project was inviting vendors to give technology demos. The demos have been a great learning experience for everyone. The vendor demonstrations will continue throughout this process. The other activity as part of the bottom-up approach was the lab work. They have provisioned a lab and built a base networking infrastructure. They have also built their first use-case scenario that they want to tackle. The idea is to plug in Zero Trust components and test capabilities using specific access scenarios.

Mr. Kerman talked specifically about Zero Trust and what it is. A traditional network with a sprinkling of subnets and requesters, generally speaking, has a perimeter placed around it. The perimeter worked for a long time until hackers became sophisticated and figured out how to penetrate the systems inside the perimeter. The traditional method had shortcomings especially when a breach happened from inside the network. The attack surface is so large there is no mechanism in place to prevent the threat vector's lateral movement. This is where Zero Trust comes in as it proposes the use of more granular perimeters that provide more control of high value assets. In the traditional environment there is implicit automatic trust. In the Zero Trust approach there is no implicit trust because you cannot differentiate between the good guys and the bad guys. Through the Zero Trust approach all data and computing searches are considered sources. All communication is secure regardless of network location.

Access to individual enterprise resources is granted on a per connection basis. A user is authorized for a specific resource but cannot use that connection for accessing another resource. Access to a resource is determined by dynamic policy, including the observable stake of user identity and the requesting system. It may include other behavioral traits as well. The enterprise ensures all owned and associated systems are in the most secure state possible and monitors systems to ensure they remain in the most secure state possible. All resource authentication is dynamic and strictly enforced before access is allowed.

There are assumptions based on where a user is, what assets a user has, and what kind of network the organization has regarding a Zero Trust view. The assumptions break out into an enterprise owned and non-enterprise owned network infrastructure. No device is inherently trusted.

The architecture receives much input from the environment including the trust algorithms. An access request, for example, is about the device. Others inputs include the operating

system version, application, and patch level. All of the information is fed into the trust algorithm. The attributes of the user, the policies, or the group come from the user database. The profile information about a specific device in the network comes from the system database. Threat vectors are also considered. There are variations for implementing trust algorithms. One variation is criteria versus score based and another is singular versus contextual. Some of the deployment scenarios include an enterprise with satellite facilities or a multi-cloud enterprise.

An event was held in November with over one-hundred and fifty attendees. Federal agencies presented on their current journeys to Zero Trust. Outcomes from the meeting provided information that requires further research and work.

The next step is a demonstration project which will be conducted at the NCCoE. The project itself has yet to be determined but could be an identity based ZTA.

A member inquired if Mr. Kerman has seen any federal agencies doing incremental implementation? And, how does that inform what NIST is doing? It has been an incremental approach. You can't just wipe out everything. A lot of agencies do not realize they are doing things that are part of ZTA which has a lot to do with lack of education on Zero Trust. One of the reasons they put out the Zero Trust document was to help people become knowledgeable about the new paradigm for securing network infrastructures.

A member asked if there any immediate near-term to two-year engineering design decisions each agency should be educated on that will take them on a path away from ZTA and should be either discouraged or assessed for risk? Not right now. Right now, it's about making use of what you have and align with existing guidance. It's an evolutionary change. We are trying to secure our environments in a better way using existing technology.

Brief from the Center for Internet Security Election Security Assistance

Ms. Phyllis Lee, Senior Director of CIS Controls, Center for Internet Security (CIS)

The Chair welcomed Phyllis Lee of the Center for Internet Security (CIS) to brief the Board on CIS's election security assistance efforts. Ms. Lee is currently Senior Director of CIS Controls; a resource CIS produces that offer a recommended set of actions for cyber defense that provide specific and actionable ways to thwart attacks. CIS has two organizations underneath it: Security Best Practices, which is where Ms. Lee works and where the work she's briefing the Board on today takes place, and the Multi-State (MS)-ISAC, which is funded through a cooperative agreement with DHS. Under the MS-ISAC is the Election Infrastructure (EI)-ISAC, which formed last year and is formally recognized as an ISAC by DHS. The work that Ms. Lee's group coordinates is with the EI-ISAC, but the work

her group does on security best practices is separate. The elections security work happens under Ms. Lee's Controls group.

The Controls are CIS's flagship document. The Top 20 Controls are a set of activities put into effect on the network to help defend it; there are 171 sub-controls under the Controls. There have been over 20,000 downloads of the Top 20 Controls document over the past four years. The Controls are created and maintained by a volunteer community, with CIS serving as the editors of the document. The group is in the process of working on Version 8, which will be released in 2021.

The election security best practice guidance is based on the Controls. CIS released the Handbook for Election Infrastructure (the Handbook) in 2018, with plans to update it in 2020. The election security guidance is intended for election officials as well as technology providers. The Handbook is loosely based on the Controls; however, the guidance related to election security contains some information not included in the Controls—88 best practices in all. The Handbook includes a security best practices guide for non-voting election technology. In 2020, they plan to release a document on supply chain guidance for elections as well as some benchmarks. The Handbook is divided into two parts: internet connected and not internet connected. In addition to the Handbook, CIS has also developed the Election Infrastructure Assessment Tool (EIAT) that enables an agency to conduct an election security self-assessment.

So far, two states that were the alleged victims of election interference, New York and Florida, created legislation mandating they evaluate their election infrastructure. CIS worked with both states to help implement the EIAT. In addition to providing the tool, CIS offers a six-step training process to help election officials work through the EIAT.

CIS also created an elections security procurement guide, funded by the Democracy Fund. The procurement guide offers example questions to ask technology providers. The guide also offers suggested language to include in RFPs, guidance on how to differentiate between good and bad responses, how certain replies are applicable, and whether responses reference the correct guidance, etc. All resources on the CIS website are free to anyone. The structure of the Handbook was originally broken up into five technology areas – software applications, servers, workstations, networking, and architecture. The volunteer committee decided to create three profiles, with Profile 1 being the minimum baseline. Profile 1 was the recommended first step, then go on to Level 2 then Level 3.

A member asked that since poll workers and election administrators aren't necessarily IT administrators, does this guidance take that into consideration? Yes, the committee was keenly aware of this fact, which is why they developed the three-level model, with Level 1 requiring only minimal technical knowledge and technology assets. Several members

asked about specific examples that went into creating the guidance—e.g. DoS, or lack of internet connectivity for non-malicious reasons. Those types of examples are incorporated into the guidance. Ransomware is currently a high-profile issue.

The latest effort Ms. Lee's team is working on is non-voting election technology verification. CIS held a workshop in November where they discussed a new process for non-voting technology verification called Rapid Architecture Based Election Technology Verification (RABET-V). Based on interest expressed in the workshop, a 2020 pilot program on RABET-V is planned. While many were enthusiastic about RABET-V at the workshop, concerns were expressed by some. A member asked whether any vendors had IP concerns and Ms. Lee replied that it wasn't that vendors were concerned about IP; rather that election officials wanted to ensure that voting machine usability and accessibility were also considered rather than just election security, because if mobility and accessibility are not considered, then some people are not physically able to vote. Another issue election officials felt strongly about was consumer confidence in the voting process in terms of making sure that every individual vote gets counted. To address these concerns, CIS is going to publish supply chain guidance, based on NIST 800-161.

The last item slated on CIS's current roadmap is elections benchmarks. The first, which is currently underway, is for Windows 10 EMS Gateway security. Benchmarks are another part of security best practice. The communities develop these as they do Controls. The Benchmarking communities are working on Windows 10, the EMS based on Windows 10 IOT, which is planned for next year, as well as Microsoft Azure for elections, which Microsoft reached out to CIS about. AWS is also interested in benchmarking AWS for Elections. A member asked whether the benchmark is like a security configuration on Windows 10, and Ms. Lee answered that it is.

A Board member asked what the remaining gaps in election security work are, as well as what worries Ms. Lee about what they're able to do and what they'd like to be able to do? Implementation is always a challenge, especially with regard to local elections officials. She hopes the existing guidance is of use to them, but they have not received any positive feedback regarding whether or not the guidance is being implemented. Big states are having to hire third-party contractors to assist with implementation, which suggests that the guidance might be too complicated for elections officials, as well as time-consuming. A challenge is simplifying what CIS does about guidance and evaluation so that it's within reach of all municipalities.

A member asked about how those states hiring contractors to assist with these efforts are vetting the capabilities of those contractors. CIS has provided training to one contractor, who was actually a former CIS employee. Another member asked whether CIS had any plans for guidance on disaster recovery or incident response and recovery in the future. It

would probably be under the purview of the EI-ISAC because incident response is part of their charter. She would check with the EI-ISAC to see if this is something they're working on or already have, as it is definitely something they should be doing in coordination with DHS.

A member asked if CIS has looked at the Cyber Independent Testing Lab (CITL) with regard to the provider process in the state development requirements, and suggested that if not, they should consider it. They probably are but she will look into it.

Overview of the DHS Election Security Guidance

Mr. Jonathan Halperin, Cybersecurity Division, Cybersecurity and Infrastructure Security Agency (CISA)
Jeff Hale, Cybersecurity Division, Cybersecurity and Infrastructure Security Agency (CISA)

The Chair welcomed Jonathan Halperin and Jeff Hale to brief the Board on DHS Election Security Guidance. Since the election infrastructure was designated as critical, the Cybersecurity and Infrastructure Security Agency (CISA), has been assisting election officials and the vendors that support them to assess and manage risks to their systems.

CISA's priority is providing assessments and resources to the field to help characterize and understand different component level risks on different aspects of how elections are administered. Business networks are one large component of and essential to the administration of elections. Once a statistically valid sample of assessment findings are found, derivative products are created because the agreement is to share the assessment findings with the party being assessed. Much of the work is determining what the commonalities are across the sector. Jonathan Halperin was the author of many pieces that have been pushed out to the community at large and will walk through some of the support he provides to election administrators.

After doing some of the assessments, similarities between the state, local, tribal, and territorial (SLTT) networks arose. Given the similarities, CISA found they could provide some nonbinding guidance that would help organizations with low budget constraints and low staff. They also prioritized areas for election administrators that would help get the most for their money as far as election security is concerned. CISA focused on low or no cost election security guidance through 'Security Tip (ST19-002): Best Practices for Security Election Systems'.

First noted as a best practice is software and patch management. One piece to such a program is the establishment of an enterprise-wide inventory list. Mitigations include implementing application whitelisting. Failure to deploy timely patch management can make an organization a target of opportunity. Automatic patch updates are available for organizations. Organizations can subscribe to the National Cybersecurity Awareness System for alerts about security updates, threats, and vulnerabilities. Another best practice

is log management. Retaining and adequately securing logs from both network devices and local hosts supports triage and remediation of cybersecurity events.

Organizations can limit the impact of cybersecurity by enforcing network segmentation. Proper network segmentation is an effective security mechanism to prevent an intruder from propagating exploits or laterally moving around a network. Another recommendation is with respect to blocking suspicious activity. Organizations should follow best practices in disabling network protocols that can be used to spread malware such as: enabling security features, scan all incoming emails, training employees to recognize phishing attempts, triaging phishing emails, and blocking macros. Credential management is becoming increasingly important. Multi-Factor Authentication (MFA) can help prevent adversaries from gaining access to an organization's assets. Establish a baseline for host and network activity. A baseline provides an organization a greater chance of finding an anomaly and responding to it. Developing and maintaining guidance and policies targeted to specific situations and assisting in implementing best practices throughout an organization benefits an organization's IT ecosystem. CISA also provides recommended elements in computing system notice and consent banners and provides an example banner. Notice and consent banners can require tailoring based on the specific circumstances and legal jurisdiction at issue. CISA also provides an abundance of additional resources to organizations on their 'Security Tip (ST19-002) site.

Before the CISA incident response teams go on site they send a questionnaire to an organization regardless of the background so that organizations or election officials can do their own self-assessment. The self-assessment provides the ability to identify areas they could improve upon.

Many election officials are also responsible for an endless amount of additional jobs on top of their election duties, which in addition to not being digitally native, makes it very difficult to ingest some of the information. CISA provides a step-wise approach to improving their security. CISA recommends that election officials, as an initial step, visit the 'Elections Infrastructure ISAC' site on the Center for Internet Security site. Additionally, CISA offers FedVTE training, 'The Election Official as IT Manager'. They offer remote penetration testing as well as risk and vulnerability analysis. CISA also provides phishing campaign assessments since phishing is a frequent attack factor. Resources are provided to help assess the maturity of an organization, which includes how to access one's web presence. Microsoft and others provide free services for both campaign and election officials to help them have better levels of security on their email and social media accounts. CISA recommends assessing the maturity of an organization's ability to rebuild their network in the case of destructive malware.

During an onsite review, CISA starts at voter registration and poll books and voting

machines. CISA has recently invested heavily in the critical product evaluation which is conducted by Idaho National Labs doing open-ended vulnerability testing. This allows them to take deployed voting systems and identify ways they can be compromised. They work with vendors to improve the machines before they go back through the Election Assistance Commission (EAC) certification process.

A Board member inquired if the list of priorities from lessons learned was derived from CISA's on-site assessments. These are areas you are recommending for improvements and not necessarily where you saw best practices? Yes, after reviewing many of the recommendations we saw many similarities across the board. A member asked how much in the election world do they see obsolete applications for loading machines, obsolete devices, or devices that depend on obsolete technology as a limitation? A lot of voting machines have yet to move to Windows 10. Across the enterprise networks, there are some unsupported applications and unsupported systems.

A Board member asked if the assessments are no cost, and if there is a menu of assessment options to choose from, why don't people/organizations choose the entire menu every time? Assessments take time and can push people out of their space for a while. There is no more valuable resource to election officials than time. A member asked what CISA is doing with the vendor community? All the services CISA provides are available to the vendor community. They are able to provide their services to any audience and election officials and vendors are critical infrastructures. As mentioned earlier, CISA provides critical product evaluation which really targets the vendor community.

Final Board Reviews and Discussions

The following areas were discussed by the Board in its review:

1. The Board will solicit themes to discuss during the March ISPAB meeting. No topics were officially proposed but one discussed was around the areas of testing (assessment conformance and compliance, etc).
2. The Board discussed potentially meeting at the NCCoE for the March 2020 meeting.
3. March 25 and 26 is the potential next meeting date. Location: TBD. The date may shift if needed.

Meeting Recessed

The meeting adjourned at 4:32 p.m., Eastern Time

List of Attendees

Last Name	First Name	Affiliation	Role
Brewer	Jeff	NIST	DFO
Badanes	Ginny	Microsoft	Presenter
Blaze	Matthew	Georgetown University	Presenter
Cofield	Val	Cyberspace Solarium Comm.	Presenter
Hale	Jeff	CISA	Presenter
Halperin	Jonathan	CISA	Presenter
Howell	Gema	NIST	Presenter
Kerman	Alper	NIST	Presenter
Lefkowitz	Naomi	NIST	Presenter
Lee	Phyllis	CIS	Presenter
Leichtman	David	Microsoft	Presenter
Romine	Charles	NIST	Presenter
Scholl	Matt	NIST	Presenter
Stine	Kevin	NIST	Presenter
Thompson	Janie	US House Comm. on Science & Tech	Presenter
Wilshusen	Greg	GAO	Presenter
McGary	Maggie	Exeter Government Services	Staff
Petrella	Evie	Exeter Government Services	Staff
Boutin	Chad	NIST	Visitor
Brown	Peter	European Parliament	Visitor
Funn	Kelby	SEC	Visitor
Geller	Eric	Politico	Visitor
Neeson	Emily	Lewis-Burke Associates	Visitor
Randhawa	Sunjeet	Broadcom	Visitor
Reeves	Terrance	NGA	Visitor
Souppaya	Murugiah	NIST	Visitor
Weber	Rick	Inside Cybersecurity	Visitor