Annual Secure and Resilient Cyber Architectures Invitational

Overview

May 2016 marked the sixth year in which more than 120 subject matter experts (SMEs) in cyber resiliency from government, industry, and academia came together in McLean, VA, for collective work on topics of common concern. For two days, the Sixth Annual Secure and Resilient Cyber Architectures Invitational (previously referred to as a workshop) accelerated recognition and adoption of cyber resiliency with a focus on organizations.

Background: 2010–2015

The first workshop, held in October 2010, established the initial community and shared architectural, technical, and policy perspectives on cyber resiliency. The second workshop, held in May 2012, focused on collaborating to develop a communal view of resiliency frameworks, engineering principles, and metrics [1]. The third workshop, held in June 2013, centered on identifying favorable conditions for use of specific resiliency techniques, assessing the use of techniques in enterprise architectures, and developing use cases [2]. The fourth meeting, now renamed "Invitational" and held in May 2014, emphasized applying cyber resiliency to space-based systems and critical infrastructure, designing a cyber resiliency challenge, and identifying roles played by cyber resiliency throughout the systems engineering life cycle [3].

The Fifth Annual Secure and Resilient Cyber Architectures Invitational, held in May 2015, concentrated on taking stock of the state of cyber resiliency: the lessons learned and the remaining challenges to overcome. It sought community consensus on the theme of *Cyber Resilience: Looking Backward (What Has Worked? What Has Not?), Looking Forward (What New Challenges Must Be Faced?)*. Keynote speakers included representatives from the National Institute of Standards and Technology (NIST), US Navy, Indiana University, and Bit9 + Carbon Black [4].

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.
©2017 The MITRE Corporation. All Rights Reserved.

2016

The most recent invitational, which took place on 18–19 May 2016, centered on the theme of *Institutionalizing Cyber Resiliency*. The invitational began with four keynote addresses by representatives of government, industry, and Federally Funded Research and Development Centers (FFRDCs). The keynote addresses were followed by an industry panel discussion and presentations by three Birds-of-a-Feather working groups. In addition, vendor booths and representatives displayed leading-edge cyber resiliency offerings. Meaningful conversations on cyber resiliency continued long after the close of the event. As one participant commented:

> *In a short period of time, thanks to the MITRE [Invitational], the topic of resilience and its relationship to 'security' has been introduced and widely socialized. The different understandings, taxonomies, and meanings for terms have been surfaced in a manner that first enabled us to see how often we are speaking past each other, and, to enable us to start closing this communication gap. I do not see how we could have come as far as we have without [the invitational].*

These proceedings present a summary of the keynote talks, the panel discussion, and the working group tracks. The Cyber Resiliency Invitational Committee hopes that in the long term the proceedings do even more: that they provide a record of growth and advancement in the influential field of cyber resiliency.

Additional materials from the invitational and briefings can be found at https://www.mitre.org/cyberworkshop The committee welcomes comments from readers through the contact email address: secureandresilient@mitre.org.

The Cyber Resiliency Invitational Committee
February 2017

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

Acknowledgments

The Cyber Resiliency Planning Committee wishes to thank the following vendors for providing booths and representatives at the 6th Invitational: *Attivo Networks, Carbon Black, CyberSponse, Cyphort, Digital Guardian, Forcepoint, RedSeal, Sonatype, and Vormetric.*

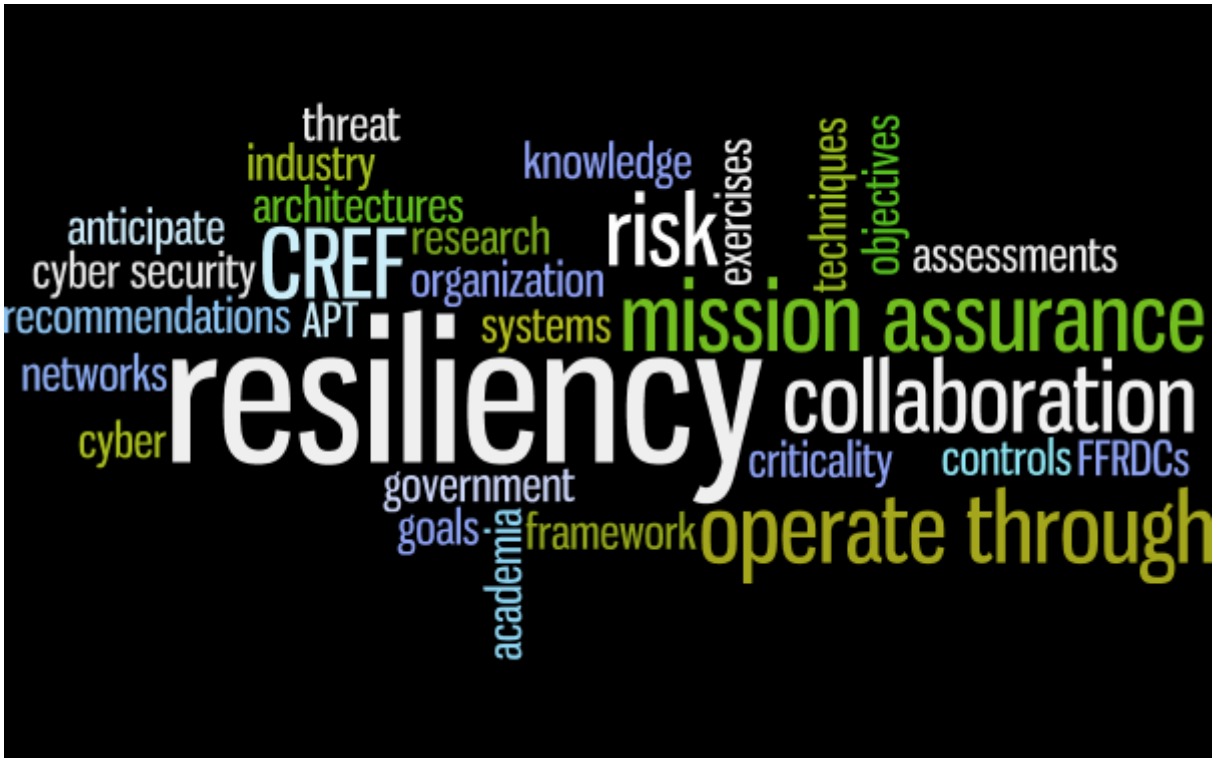

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

This page intentionally left blank.

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

Table of Contents

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

List of Figures

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

1. Introduction

The Sixth Annual Secure and Resilient Cyber Architectures Invitational brought the cyber resiliency community together to explore the impact of cyber resiliency on organizations. The invitational also examined the effects that organization characteristics have on efforts to include cyber resiliency in cyber security programs. Four keynote addresses, one panel, and three facilitated working groups provided the structure and content to tackle 2016's far-reaching theme.

Section 2 summarizes the four keynote addresses and one panel, as follows:

- "Use of Deception to Foster Resilience," given by Kristin Heckman and John Woodward, The MITRE Corporation
- "Resiliency and the CISO," given by Arlette Hart, Chief Information Security Officer (CISO), Federal Bureau of Investigation (FBI)
- "Is Cyber Resiliency Really that Difficult," given by John Gilligan, President, The Gilligan Group
- "Technical and Usability Factors: Validation of Resiliency," given by Nick Multari, Principal Program Manager, Cybersecurity Research, Pacific Northwest National Laboratory (PNNL)
- "Industry Perspective on Cyber Resiliency," chaired by Shane Steiger, Hewlett Packard Enterprise (HPE).

Section 3 provides details on the three working groups, as follows:

- Cyber Resiliency and an Organization's Cyber Security Program
- Cyber Resiliency and Architectural and Engineering Processes
- Cyber Resiliency in Acquisitions.

As a method to reach out and broaden the technical community engaged in cyber resiliency, MITRE provided a three-hour tutorial session directed at those participants relatively new to the field. The agenda for the tutorial, held the day before the invitational, focused on cyber resiliency in four foundational ways:

- An overview covering motivations, definitions, relationships to other disciplines, and introduction to MITRE's Cyber Resiliency Engineering Framework (CREF)
- Cyber resiliency analyses and their end products
- Cyber resiliency and metrics: descriptions, challenges, and examples
- Cyber resiliency and the Risk Management Framework (RMF).

MITRE hopes to continue this practice as more – and different – types of organizations expand their cyber security approaches to encompass resiliency.

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

2. Keynotes and Panel

2.1 Use of Deception to Foster Resilience, *Kristin Heckman, John Woodward, The MITRE Corporation*

The keynote began by reviewing a typical scenario in cyber resiliency, i.e., fighting through an attack. It was noted in recent years this typical scenario has changed and expanded, especially regarding high-end offenses, whose most likely targets might be closed networks, such as industrial controlled networks and classified environments. Such offenders are likely to steal information to be used to change their own (adversarial) behavior, thus making attacks more formidable.

Consequently, deception is one of the most effective cyber resiliency techniques available. The keynote continued by presenting the goal of all deception effects: namely, to influence an adversary's action (or inaction) and then to benefit from those decisions. Deception is one of the fourteen techniques specified in the MITRE CREF.

"Denial actively prevents the target from perceiving information and stimuli; deception provides misleading information and stimuli to actively create and reinforce the target's perceptions, cognitions, and beliefs." [5] The speakers presented cyber denial and deception (D&D) types and tactics as well as a methods matrix. Types and tactics cover both *show the real* (e.g., paltering, negative spin, feints/demonstrations, double play, and double bluff), and *hide the real* (e.g., masking, repackaging, dazzling, and red flagging). Likewise, types and tactics can cover *show the false* (e.g., mimicking, inventing, decoying, and double play), and *hide the false* (e.g., operational security (OPSEC) and positive spin). As part of the presentation, Ms. Heckman introduced the audience to the concept of a deception chain, as shown in Figure 1. The deception chain is an analogous model to Lockheed Martin's "cyber kill chain" model [6]. It works to integrate three systems – cyber D&D, cyber intelligence, and security operations – into an organization's larger system of deception operations.

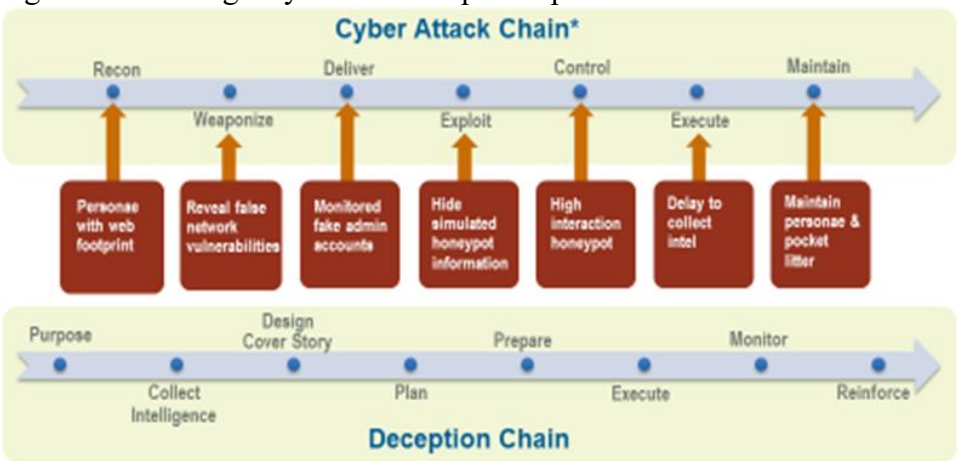

Figure 1. Cyber Attack Chain, Deception Chain

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

The speakers also gave a walk-through of cyber D&D used to thwart a realistic advanced persistent threat (APT). They described the benefits of cyber D&D in detail; these benefits include increased insight into the techniques and intents of an adversary, an improved ability to tailor defensive and offensive approaches to threats, the strengths of an evolving incident response, and overall enhanced intelligence about the adversary.

2.2 Resiliency and the CISO, *Arlette Hart, CISO, FBI*

Arlette Hart opened the second keynote address by drawing distinctions between cyber security and cyber resiliency. She explained broad implications of the historical, technological shift from single-purpose devices with physically located-dependent data to the present-day Internet-of-Things (IoT). Ms. Hart showed that intrusions at the workplace may now come from homes and elsewhere because of the proliferation of multi-purpose devices: think of household refrigerators with door displays of family schedules pulled from on-line calendars. Workplaces may end up with "inadvertent insiders" permitted from a lack of proper security on IoT devices. In many fields, the increased use of software as a service (SaaS) has taken the management and configuration of workplace systems out of the direct control of the organizations being served. Furthermore, Ms. Hart asked participants to consider that the security of critical infrastructure may depend on a very thin technological margin, with no in-depth defense of mission-critical functions.

The IoT can contain consumer items that generate unintended consequences. Ms. Hart cited examples from the daily news such as toy Furbies that record environmental sound as well as a child's voice. Televisions with on-board cameras may record the actions and voices of viewers without explicit permission, thus technologically permitting illegal eavesdropping. Ms. Hart also brought to light that over the past two decades both government and commercial systems have been built with software developed to meet the primary objectives of speed and efficiency, often with little regard for cyber security. Therefore, a broad, non-secured base of software exists with the end results of expanding the attack surface.

Ms. Hart presented resiliency activities and cybersecurity controls in the context of society at large. She discussed the perspectives of CISOs of larger organizations, which included identifying organization-owned data, building understanding of the impact of possible data breaches and compromised capabilities, and determining adversaries and their possible intents. CISO priorities include identification of data assets, knowledge of supporting capabilities, impact of asset degradation or loss, and integration of resiliency with an organization's security program. Ms. Hart stressed the importance of including cyber resiliency in the language of all businesses, as well as the importance of incorporating cyber resiliency into business operations across industries. She ended her talk by reviewing incident reports and response plans, and noting how concepts such as coordinated defense and segmentation – both cyber resiliency techniques specified in the MITRE CREF – can help mitigate threats.

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

2.3 Is Cyber Resiliency Really That Difficult? *John Gilligan, Center for Internet Security (CIS) and President, The Gilligan Group*

John Gilligan began his keynote address by framing the talk from his personal journey in cyber resiliency. He realized that his original dream of resiliency represents a very long-term objective because of the very nature of cyber resiliency: a complex, system of systems engineering challenge. He also highlighted that cyber risk management requires knowledge that most organizations do not yet possess in house and that organizations face market forces poorly aligned with achieving resiliency.

However, achieving high resiliency is possible today through a structured journey, optimally with cyber resiliency well integrated into an organization's cyber security plan. Mr. Gilligan reviewed a strong framework, a top-level resiliency strategy, and implementation steps. At the foundation of each was Version 6 of the CIS comprehensive baseline of security controls. Of particular note were controls 1–5, circled in Figure 2. Mr. Gilligan pointed out that these controls are referred to as "good security hygiene." The Australian Signals Directorate Study claimed that more than 80 percent of cyber threats might be avoided due to successful adoption of such controls.[1]

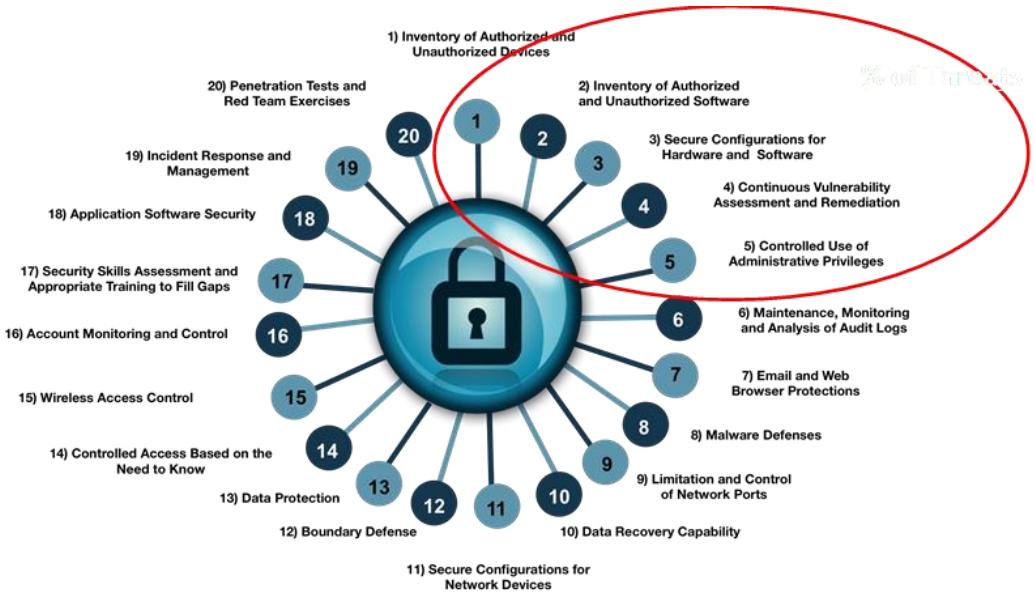

Figure 2. CIS Comprehensive Baseline of Security Controls

Mr. Gilligan presented a framework for cyber resiliency maturity in which he explored five levels, as shown in Figure 3. He based the discussion on Robert Lenz's maturity model [7]. Level 1, No Resilience, is characterized by inconsistent deployment of security controls. Mr. Gilligan speculated that many of today's largest organizations are at Level 1. Level 2,

1 https://www.asd.gov.au/infosec/mitigationstrategies

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

Performed, involves the implementation of foundational/critical security controls (CSCs). Level 3, Managed, is reached when CSC integrated controls are in place and continuously monitored. (Levels 1–3 are considered Step 1 in establishing resiliency and referred to as the CSC Baseline.) Levels 4 and 5, Dynamic and Resilient, assume an augmented CSC baseline based on mission information and needs. Level 4's threat response is rapid reaction with responses to sophisticated cyberattacks. Organizations operating at Level 5 exhibit anticipation of threats and can operate through sophisticated attacks. Mr. Gilligan concluded the talk with an emphasis on taking the long view to achieve cyber resiliency.

Cybersecurity Resilience Maturity Framework

	Maturity Descriptor	Employment of Security Controls	Security Tailored to Mission	Participate in Information Sharing (threat/vul.)	Response to Cyber Threats	Resilience to Cyber Attacks
Step 2: Address Sophisticated Attacks	Level 5: Resilient	Augment CSC Based on Mission	Mission Assurance Focused	Real Time Response to Inputs	Anticipate Threats	Operate Through Sophisticated Attack
	Level 4: Dynamic	Augment CSC Based on Mission	Mission Focused	Real Time Response to Inputs	Rapid Reaction To Threats	Able to respond to Sophisticated Attack
	Level 3: Managed	CSC Integrated and Continuously Monitored	Partially Mission Focused	Respond to Information Inputs	Respond to Attacks After the Fact	Protection against Unsophisticated Attack
Step 1: Implement CSC Baseline	Level 2: Performed	Foundational/ Critical Security Controls (CSC) Implemented	Mission Agnostic	Inconsistent Response to Information Inputs	Respond to Attacks After the Fact	Some Protection Against Unsophisticated Attacks
	Level 1: No Resilience	Inconsistent Deployment of Security Controls	None	None	No Response	Susceptible to Unsophisticated Attacks

Figure 3. Cybersecurity Resilience Maturity Framework

2.4 Technical and Usability Factors: Validation of Resiliency, *Nick Multari, Principal Program Manager, Cybersecurity Research, Pacific Northwest National Laboratory (PNNL)*

The mission of PNNL's Asymmetric Resilient Cybersecurity (ARC) Initiative is to "deliver the theory, processes, methodologies, and algorithms that will enable a resilient cyber infrastructure with an asymmetric advantage, thwarting adversaries who seek to infiltrate and damage our national security through digital means."[2]

2 http://cybersecurity.pnnl.gov/

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

The three components of validation are theory, technology, and usability. Each is critical to the successful institutionalizing of resilience technologies. Mathematics and logic proofs establish the theoretical foundations. The usefulness of theories is then verified in practice by applying theoretical concepts to ensure critical missions can withstand an attack. The validation process comprises first proof-of-concept demonstrations followed by demonstrations of capabilities in specific environments.

As ARC develops a technology validation plan, team members build three scenarios: destroy, steal, and damage. Background traffic is simulated using Lincoln Laboratories' LARIAT.[3] For demonstrations, ARC uses the entire testbed environment with full integration of all theories and technologies.

Currently, while tabletop exercises exist that address resilient infrastructure, none appear to evaluate human usability factors (i.e., those incurred when implementing resilient technologies). Future partnerships to explore this area would help to ensure organizations' acceptance of resilient infrastructures.

The talk concluded by describing four studies that explored cybersecurity decisions and the cost of such defenses. Participants in the four studies came from academia, industry, government, and research organizations. The studies all followed a tabletop exercise model with four participants per study, each assuming the role of security, business, engineering, and intelligence, respectively. Participants were given technologies and resources to build their networks, with Year 1 of the studies using non-resilient technologies and Year 2 using resilient technologies. Initial results from the studies showed that all teams adopted similar initial approaches. In addition, all teams identified some investment for which they desired to obtain resources later. However, differences were observed in participants' choices of strategies concerning the breadth of technologies (low vs. high level), depth of technologies (few, but high level), or a combination of the two.

2.5 Panel: Industry Perspective on Cyber Resiliency

Chair: Shane Steiger, Hewlett Packard Enterprise (HPE)

The panel discussion opened with three use cases of major data breaches described in the media over the past two years: Target©, Sony©, and Medstar Health©. As part of its discussion, the panel pointed out that the lack of applicable cyber resiliency techniques may have facilitated the attacks.

The discussion moved on to thirteen industry-developed cyber resiliency guidance areas (listed below) related to the lifecycle of a cyberattack that were developed for executives and system

3 LARIAT (Lincoln Adaptable Real-time Information Assurance Testbed) is capable of emulating networks consisting of one to one million physical hosts, and modeling users performing real tasks, with real application software, whether checking e-mail and browsing the web or operating military sensors and weapon systems.

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

architects. Six of the guidance areas address planning and preparation activities, six cover recovery and reconstitution activities, and one applies to both. Specifically, these areas were:

- Planning and Preparation Activities (Before Boom)
    - Architect to Protect
    - Secure Administration
    - Access Control
    - Device Hardening
    - Backup Strategies
    - Cyber Continuity of Operations (COOP) Planning
- Recovery and Reconstitution Activities (After Boom)
    - Cyber COOP Execution
    - Secure Communications
    - Core Services
    - Data Recovery Strategies
    - Forensics
    - After Action Report
- Disrupting the Attack Surface – Overarching Activity

Readers interested in additional details on industry-developed cyber resiliency guidance are directed to content from the invitational [8].

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

3. Working Groups

3.1 Cyber Resiliency and an Organization's Cyber Security Program

Lead: William Knox, Harvard University
Deputy Lead: Ellen Laderman, The MITRE Corporation

3.1.1 Goals

This working group had two goals. First, the group members addressed identification of challenges to incorporating cyber resiliency into an organization's cyber program. Second, they offered recommended guidance for cyber professionals with the additional goal of lowering the barrier to acceptance of resiliency in organizations.

3.1.2 Discussions/Observations

The discussion and observations revolved around three themes: the context of missions and their associated threats, the cultures of organizations, and complexities of environments and their requirements. Track participants discussed problems and potential solutions for each of these themes.

The concepts of both mission and threat vary depending on the context in which one examines them. Consider, for example, the case of an unmanned aerial vehicle (UAV), when the mission of an individual UAV is likely part of a larger operational plan. Understanding threats to the mission also requires context, but gaining this understanding is difficult because organizations frequently do not have access to classified threat information. Furthermore, if an organization has access to threat information, it must still ask, "What is the adversary's intent once entry is obtained?"

The solutions to these challenges range from activities within the organization to changes in the broader environment. Intra-organization activities include creating and maintaining inventories of software and hardware, and using publicly available threat information such as MITRE's ATT&CK [Adversarial Tactics, Techniques, and Common Knowledge] and the Information Technology –Information Sharing and Analysis Center (IT-ISAC).[4] Organizations must also consider those activities that may require help from outside contractors who, perhaps inadvertently, gain access to sensitive threat information. For example, organizations should carefully monitor heating, ventilation, and air conditioning (HVAC) service contractors who have blueprints and physical access.

4 https://attack.mitre.org/wiki/Main_Page and http://www.it-isac.org/

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

Evaluation of risk acceptance must occur in the context of how well and how quickly an organization could respond to risk once it materializes. Organizations need to know the context for both the mission and the threat in order to evaluate risk in a rational manner. While the organization is responsible for seeking help and advice, the standards organizations and government must inform them and provide such help.

An organization's culture may present one of the biggest challenges to incorporating cyber resilience into the organization's cyber program. This can be summarized by the quotation "culture eats strategy for breakfast."[5] No matter how superb an organization's strategy is, it will fail if it is not supported by the culture. Track participants presented four illustrative examples of cultural hindrances.

First, one challenging cultural belief is that if security fails *anywhere*, then it has failed across the board (e.g., a vulnerability in one system is a vulnerability shared by all). An additional commonly seen and experienced barrier is called the "check the box" culture: one dominated by security checklists with an overwhelming number of items, which is counterproductive to the successful application of controls. Third, accreditors and data stewards may erroneously use cyber resiliency as the sole standard. The group members recommended that cyber resiliency be fully incorporated into an organization's current cyber security standards; otherwise, an organization runs the risk of gaps and omissions in coverage, with unmet security standards for a given mission. The fourth and final barrier discussed was the possible omission of gathering input from all stakeholders in an organization. The extent and complexity of this barrier vary directly with the size and complexity of the organization itself.

The track participants discussed two areas of solutions at length, one concerning advocacy efforts and another concerning process change.

Advocacy efforts entail empowering resilience experts to champion cyber resilience methodology and requirements. Such advocacy efforts can succeed in breaking down the above-mentioned cultural barriers. The group offered the practical suggestion to phrase cyber resiliency needs in terms of an organization's existing government/military culture, for example, refer to "mission assurance" in place of "cyber resiliency," if appropriate.

The working group defined process change as a sequence of steps or activities that a team or project leader follow to drive individual transitions and ensure the project meets its intended outcomes to complete its mission.[6] Regarding cyber resiliency, the group offered four suggestions for successful process change. One suggestion was to adopt the Structured Cyber Resiliency Analysis Methodology to better communicate the relationship between cyber security and cyber resiliency.[7] This methodology provides prioritized recommendations rather than lists, enabling an organization to stop treating baseline controls as mere checklists and start integrating

5 This quotation has been attributed to Peter Drucker http://www.forbes.com/sites/shephyken/2015/12/05/drucker-said-culture-eats-strategy-for-breakfast-and-enterprise-rent-a-car-proves-it/#6b0a8ad574e0
6 https://www.prosci.com/change-management/thought-leadership-library/change-management-process
7 https://www.mitre.org/sites/default/files/publications/pr-16-0777-structured-cyber-resiliency-analysis-methodology-overview.pdf

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

risk management. Another suggestion was to use Appendix H, System Resiliency, of NIST SP 800-160, *System Security Engineering*, to improve existing compliance cultures by incorporating cyber resiliency.[8] The group also discussed involving more stakeholders in tabletop exercises as a way of emphasizing the urgent need for cyber resiliency.

One of the many reasons why complex environments challenge organizations to incorporate cyber resiliency is that the task of mapping asset dependencies becomes too broad in scope. Organizations commonly have multiple missions that affect each other, with managers possibly unaware of the full mission impact. Part of the solution to the broad task of dependency mapping is to begin with mission-critical functions and then proceed through a function hierarchy, as time and resources permit.

A central portion of any solution relies on intra-organization communication and prioritization of objectives. Basic lists of actions, or tasks, are less helpful than ranked recommendations based on cost-benefit analyses. In addition, a reference architecture built by a "coalition of the willing" may serve as a starting point for organizations to begin the involved task of incorporating cyber resiliency. Such a reference architecture of a resilient environment would also include the architecture's components (e.g., directory, storage/cloud, key store). The group expressed the hope that providing a detailed model would help more organizations to embrace process change.

3.1.3 Challenges

The track participants identified four challenges to incorporating cyber resiliency. Three of them corresponded to organizations' cultures, complexity of systems, and context of environments. The fourth challenge focused on the relative maturity of cyber resiliency techniques, thus having wide-ranging impact across many different types of organizations. Specifically, some cyber resiliency solutions are not yet fully developed. One participant referred to this as "not ready for prime time." Of course, what defines "prime time" varies from organization to organization as well as from mission to mission. In general, however, the more critical the environment is, the more stable and mature the technology portion of the resiliency solution must be.

3.1.4 Recommendations/Way Forward

The track participants developed ten recommendations to address the above discussions on challenges. In no prioritized order, the recommendations were:

- Develop cyber resiliency tabletop exercises. In addition, develop example exercises to be done before tabletop exercises. Example exercises tailor the needs of an organization, thus making tabletop exercises more effective. Both example and tabletop exercises: a) provide a point of entry to understand the environment, b)

8 http://csrc.nist.gov/publications/drafts/800-160/sp800_160_second-draft.pdf

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

highlight areas for application of cyber resiliency techniques, and c) provide ammunition for "Cyber Resiliency Heroes" to champion programs going forward.

- Incorporate cyber resiliency tabletop exercises into organizational processes. Tabletops can be used to change culture by giving decision makers realistic experiences of threats and exposures, helping them to understand the interdependencies of an organization, and testing new implementations before they become operational.

- Integrate basic cyber *resiliency* guidance into cyber *security* guidance, e.g., CSC. This would serve to align an organization's strategy to its culture. As resiliency guidance becomes more widely available, the push to integrate cyber resiliency should be encouraged.

- Develop guidance based on grouping techniques according to commonalities. For example, techniques could be grouped according to environments (e.g., embedded or cloud) or according to focus of mission (e.g., privacy or availability). An interesting grouping might address the intent of the adversary (e.g., denial of service, destruction, or exfiltration). Once again, track participants mentioned the use of the ATT&CK matrix as a way to map out techniques.

- Develop a set of architecture examples that incorporate cyber resiliency. Such examples would aid organizations in integrating cyber resiliency into their systems.

- Develop overlays for compliance- and regulatory-oriented cultures. Such overlays should highlight resiliency techniques and mitigations.

- Develop knowledge centered on the impact of cyber resiliency techniques on other mission requirements (e.g. performance and operations). This body of knowledge would place cyber resiliency in the context of operational needs and constraints.

- Attempt to influence industry to further develop a subset of technologies and map them to specific Technical Readiness Levels (TRLs). This would constitute an important first step in addressing the challenges posed by the previously mentioned "not ready for prime time" techniques.

- Change language in the field to limit the distinction between cyber *resiliency* and cyber *security*. Once stakeholders learn and understand that continuum and the overlap between cyber *resiliency* and cyber *security*, they may more willingly shift their frames of reference.

- Promote the field of cyber resiliency through education and communication. Cyber resiliency addresses many of today's emerging risks. Given the ever-growing body of cyber resiliency knowledge, organizations will become successful in incorporating such knowledge as the field is better championed.

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

3.2 Cyber Resiliency and Architectural and Engineering Processes

Lead: Dr. Ron Ross, Fellow, National Institute of Standards and Technology (NIST)
Deputy Leads: Deb Bodeau and Richard Graubart, The MITRE Corporation

3.2.1 Goal

The goal of this track was to develop guidance to advance the incorporation of cyber resiliency into the system security engineering (SSE) process.

3.2.2 Discussions/Observations

The track began with a briefing and some background discussion. The second public release of NIST SP 800-160 occurred one month before the invitational and it included Appendix H, System Resiliency. The background discussion explored the contents of key sections of the document, notably Chapter 3, The Processes, and Appendix H, System Resiliency. The track participants explained Figure 4 and provided specific guidance by describing the relationship between NIST SP 800-160 and the overall systems engineering process.

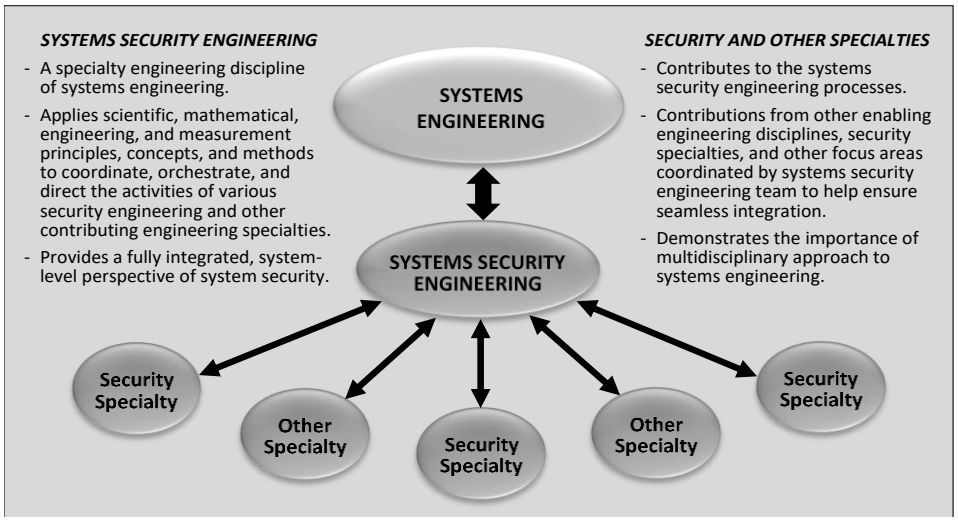

Figure 4. Systems Engineering and Other Specialty Engineering Systems

As the working group progressed, it discussed differences in emphasis and focus between NIST SP 800-160 and established cyber resiliency thinking. For example, cyber resiliency starts with the assumption that an adversary has achieved a persistent foothold within the system. This point is not explicitly noted in NIST SP 800-160. In addition, group members noted a difference in

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

goals and objectives. Goals and objectives as stated in NIST SP 800-160 include *preservation of confidentiality, integrity and availability*, as well as the need to *identify, protect, detect, respond, and recover.* By contrast, cyber resiliency goals are to: *anticipate, withstand, recover, and evolve.* Cyber resiliency objectives are to: *understand, prepare, prevent, continue, contain, reconstitute, transform and re-architect.* Such terminology is represented in MITRE's CREF, as shown in Figure 5.

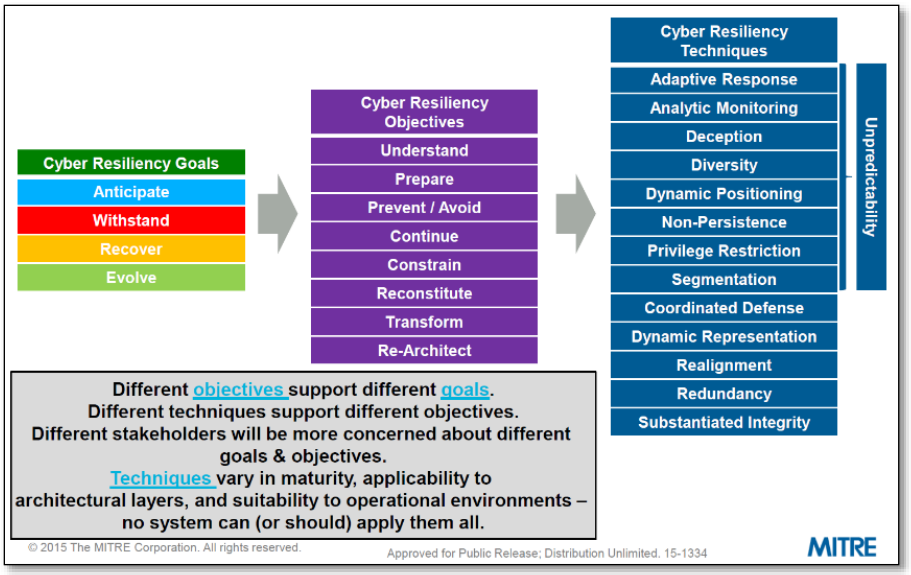

Figure 5. Cyber Resiliency Engineering Framework (CREF)

When defining requirements, decisions on architecture and priorities are commonly traced back to underlying objectives and goals. Since it appears that NIST SP 800-160 does not yet definitively incorporate resiliency goals, the group raised some concern regarding the impetus for decision makers to include such goals when writing future requirements.

Finally, the working group discussed the philosophy of protection (PoP) at length. PoP is a key construct in NIST SP 800-160. To be more complete, the description of PoP in the document should reflect the possibility of an adversary's successful compromise of resources. Content about PoP should also address decisions about cyber resiliency trade space, for instance, how much emphasis to place on trying to keep an adversary out versus combatting an adversary that has already achieved a foothold.

3.2.3 Challenges

A key challenge concerned standardizing cyber resiliency terminology and ideas in NIST SP 800-160 without making structural changes to the document. Furthermore, SSE concepts such as trustworthiness and software assurance should be considered integral to any writings on cyber resiliency. The participants did agree that NIST SP 800-160 cannot be changed to reflect the

concerns of all related disciplines; to do so would risk failure to meet the original intent of this influential document and, in fact, might preclude its widespread use.

That said, the group stressed the importance of aligning and reinforcing cyber resiliency and other disciplines in the overarching SSE loop. Each discipline has important perspectives and value that must not be lost. The group also understood the importance of not limiting discussion of special topics, including cyber resiliency, to a single section of NIST SP 800-160. Doing so would impede the broad usefulness of the main body of the document, and in this instance, its applicability to cyber resiliency.

3.2.4 Recommendations/Way Forward

The group agreed on the importance of connecting cyber resiliency (and other emergent disciplines) with the up-front material of NIST SP 800-160. For example, the up-front engineering guidance should include basic concepts such as resiliency and the notion of an advanced adversary gaining a foothold. It was also suggested that some key elements of cyber resiliency as presented in Appendix H, System Resiliency, should be incorporated into the main body of NIST SP 800-160, especially into Chapter 3, The Processes.

Group members also recognized that cyber resiliency and other emergent disciplines are likely to have goals and objectives that diverge from (or perhaps go beyond) traditional goals and objectives of cyber security, e.g., confidentiality, integrity, and availability. The group emphasized the impracticality of trying to include *all* such goals and objectives in NIST SP 800-160 (especially in Chapter 3). As a result, the group concluded it may be important to "well-scrub" the front of the document to remove all references to specific goals or objectives.

Additionally, the working group distinguished the relative importance of those resiliency techniques that are (or should be) standards of good practice from those whose selection is more risk driven. Some techniques, such as segmentation and privilege restriction, should be considered foundational and therefore reflected in the body of NIST SP 800-160. Conversely, other techniques, such as deception, dynamic positioning, and non-persistence, would more likely be selected based on environment specifics and each organization's risk tolerance, and as such, should not be covered in the body of the document.

The group agreed that it would be ideal to augment Appendix H, System Resiliency, with various vignettes, scenarios, or use cases. Members believed that use cases are effective in providing specific guidance on incorporating cyber resiliency into various SSE processes. The group noted that four factors should be considered in constructing such cases: 1) type of system, 2) environmental assumptions, 3) risk management strategy, and 4) lifecycle stage.

The group spent considerable time developing these four factors as follows.

Types of systems might include enterprise IT, embedded, critical infrastructure system, and system of systems. Each of these might, in turn, lend itself to more than one resiliency technique.

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

For example, techniques such as deception or dynamic positions would not be appropriate for an embedded system, whereas segmentation, privilege restriction, and non-persistence might be.

Environmental factors would span the context in which the system operates or will operate. Six sub-factors discussed were:

- Size, transaction volume, scalability – COOP and disaster recovery requirements imply a need for resiliency requirements
- External connectivity and external relationships
- Relationship of a system to a multi-tiered risk management approach – consider the role of the system in a mission or enterprise architecture
- Systems development, maintenance environment – how do cyber defender DevOps personnel fit in?
- Technical and operational environments – where can controls be allocated, what is the existing infrastructure?
- Time considerations – real-time reaction, or near-real-time reaction.

On this last point, the group gave an illustrative example. For systems requiring critical, real-time responsiveness, e.g., embedded heart monitors or guidance systems on aircraft, periodic refreshing of software, a common resiliency technique for non-persistent attacks, may not be adequate.

Risk management factors cover a broad aspect of cyber resiliency. The working group reviewed differing perspectives. For example, one organization may be more likely to focus on the early stages of an attack life cycle, trying to minimize an adversary's presence in the system, whereas another organization may use an adversary's presence to discern facts about the intruder before expunging it.

A system's place in the system development life cycle has an impact on decisions about cyber resiliency security engineering. Organizations whose systems are in the very early, developmental, stages of their life cycle may have the luxury of examining a broader suite of resiliency techniques. By contrast, a mature system will likely have a well-established architecture. Organizations with such systems will probably be more restrictive regarding the timing and nature of cyber resiliency augmentations (e.g., perhaps only during scheduled upgrades for mitigation add-ons).

The working group developed a *representative use case* for the Building Control System (BCS) in a fictional, but realistic, smart building. The group noted that the Architectural Definition & Design Definition processes should highlight relevant resiliency techniques to mitigate risks stemming from compromised field programmable gate arrays (FPGAs). Five techniques discussed were:

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

- Substantiated integrity – apply the Byzantine Quorum approach
- Privilege restriction
- Segmentation – ensure physical access control systems are logically isolated from external access (including external maintenance)
- Analytic monitoring for abnormal behavior
- Diversity and realignment (e.g., exclude FPGAs from a targeted set of technologies in the BCS architecture, or acquire FPGAs from multiple sources)

3.3 Cyber Resiliency in Acquisitions

Lead: Jeff Stanley,[9] Associate Deputy Assistant Secretary of the Air Force for Science, Technology and Engineering, Office of the Assistant Secretary of the Air Force [Acquisition])
Deputy Lead*:* Dan Holtzman, The MITRE Corporation

3.3.1 Goals

This track had two goals. The primary goal was to develop a roadmap that would aid professionals in incorporating cyber resiliency into the acquisition process. The secondary goal was to differentiate cyber resiliency from system resiliency and, more broadly, systems engineering.

3.3.2 Discussion/Observations

The main objective for the track was to plan a series of technical exchange meetings (TEMs) to occur over the next year, with a final TEM occurring at the 2017 Resiliency Workshop. During these TEMs, the group will meet to discuss important topics to be determined during this track so as to move toward better cyber-resilient weapon systems.

The track focused on systems requirements. Cyber resiliency requirements should be treated no differently than other system capability, performance, and effectiveness requirements. The requirements are:

- Derived from statements of need and associated concerns
- Optimized, balanced, and traded across stakeholders
- Stated as *system* requirements in all relevant contexts and across a variety of acquisition and engineering artifacts.

The statement of need for cyber resiliency in a Request for Proposal (RFP) should be clearly written, especially the objectives and measures of success. The statement of need must be

[9] Jeff Stanley was the scheduled track lead. He was replaced by his Deputy, Lt Col Rick Day.

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

elicited, analyzed, and assessed across all relevant contexts. The group noted that the Department of Defense (DoD) has many definitions of "cyber resiliency," for example, in the US Air Force (USAF) and the Office of the Deputy Assistant Secretary of Defense for Systems Engineering (DASD(SE)). For purposes of the track, the working group members chose to speak of cyber resiliency in terms of mission assurance because addressing resiliency is key to overall mission assurance.

Furthermore, mission assurance has many dimensions and views; no single stakeholder looks across the acquisition systems engineering landscape. Mission assurance comprises a complex combination of operational concerns (planning and execution), individual systems, inherent vulnerabilities, external factors, adversary tactics, and a set of complex interdependencies. Affordable, effective, and efficient mission assurance requires a wide and deep view across the numerous systems that support a mission with an end goal of creating an integrated capability. *Important to note, resiliency is not a specific characteristic that one mission or system acquires, but is rather an overall characteristic derived from the attributes of the capabilities and assets that support the mission.*

After discussing goals and context, the group switched to an Air Force briefing on cyber-resilient weapons systems. The briefing listed seven USAF levels of assurance (LOAs):

1. Mission thread analysis
2. Systems engineering and placing resilience into the acquisition process
3. Cyber workforce skill set
4. Agility and adaptability from a system designed in a modular way
5. Common security environment with an aggregation of vulnerability for data at risk
6. Assessment and repair of fielded systems with "bolted on" resilience based on system risk
7. Intelligence

During the briefing, the group raised several areas of concern. First, the test and evaluation community must evolve to include cyber resiliency as part of testing. Second, the group agreed on the clear need for more cyber workforce involvement. Third, questions arose regarding how cyber resiliency should be measured and what metrics should be used. The measurement of cyber resiliency will be the topic of a future TEM.

It was clear to the participants of the workshop that in acquisitions, cyber resiliency requires a view across the entire mission thread to consider: a) all the functions of the mission, and b) critical nodes that need to be hardened and, therefore, made more resilient.

Once the Air Force briefing concluded, the discussion refocused on the main goal of the track: to develop a roadmap for professionals in incorporating cyber resiliency into the acquisition process. Four main observations resulted, described as follows.

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

1. National Security Agency (NSA) Information Assurance Capabilities is working with the DoD Chief Information Officer (CIO) to develop a Cyber Survivability Endorsement (CSE) that in three volumes describes ten areas for cyber resiliency to address. The guidance in Volume 1 of the CSE centers on weapons systems. It describes how to design requirements so the system can perform its mission while incorporating cyber resiliency. The purpose is to give specific guidance to Project Managers, who may be more focused on schedule and cost and, perhaps, less familiar with resiliency. Volume 2 gives exemplars that could be directly used in RFPs. Volume 3 provides specific, classified guidance.

   To implement the CSE in practice, organizations must consider the relevant policy that addresses cyber resiliency. Four action items resulted from the discussion: to determine what policy currently exists, the gaps in policy, the incentives for Project Managers to implement resiliency, and the likely flows from policies to DoD Directives (or Instructions).

2. The group discussed "Analysis and Validation from a Whole Life-Cycle View," and concluded that such a view requires a deeper understanding of the system being purchased with an eye toward those that incorporate resiliency *and* operator tasks.

3. Increased use of commercial off-the-shelf (COTS) products has prompted an increased focus on supply chain risk management (SCRM). The Air Force is working with the National Defense Industrial Association (NDIA) and the DASD(SE) office to improve its performance with regard to SCRM, but the group acknowledged that SCRM presents a sizable and difficult problem to solve. Dependency maps to better understand the mission criticality of components are needed, but unfortunately such maps are not prepared on a regular basis at this point. The top-level goal is not to make a system perfect but, rather, to make it difficult for an adversary to get into a system and remain undetected.

4. The group also discussed the difficult-to-achieve but important goal of "Owning the Technical Baseline (OTB)." Speakers mentioned that shortfalls may occur in accomplishing the goal if the government ends up buying data rights on missed deliverables or must reverse-engineer deliverables to reconstitute a technical baseline.

In closing, the group identified additional areas of exploration that will be addressed in future TEMs: science and technology investments, OPSEC considerations, system and mission assurance practices, and requirements specification.

3.3.3 Challenges

The group brought to light five barriers to achieving cyber-resilient acquisitions: changes in the requirements process, organizational culture, financial considerations, acquisition intelligence, and the acquisition method of Lowest Price Technically Acceptable (LPTA) contracts.

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

One area of needed improvement concerned the writing of resiliency requirements in RFPs. Speakers noted that resiliency changes must be improved, from the Joint Requirements Oversight Council (JROC) all the way down to specifications and contracts. While the JROC understands the need to write better requirements for resiliency, it currently lacks a sufficient workforce to do so. One possible solution may be to provide outside support to the JROC from the technical community.

As previously noted, changing the cyber culture represents a significant challenge from many perspectives. Some organizations place their emphasis solely on cyber security compliance; others have a fundamental "keep them out" philosophy. Only a few have the fully realized cyber resiliency philosophy of operating through an attack. The group suggested that acquisition should prioritize resiliency over performance and champion a cyber resiliency philosophy.

Writing requirements depends strongly on timely intelligence as part of the acquisition cycle.

The working group also discussed a major challenge that concerned the costs of cyber resiliency. Speakers noted that in some cases, if the program has no hard requirement for resiliency, that feature may be removed during budget cuts. Additionally, because COTS vendors are driven by cost, they have no financial incentive to include resiliency in government projects when the vendor's commercial marketplace does not yet require it. One line of reasoning suggested that if the government were to require a trusted chain of suppliers to attack the existing SCRM problem, this would lead to higher initial costs in system acquisitions but certainly produce savings further down the road.

At the end of the working group session, speakers raised the challenge presented by LPTA. The group concluded that LPTA can leave resiliency behind and out of the process. A future solution may involve using a mission assurance cost effectiveness metric instead.

3.3.4 Recommendations/Way Forward

The working group made recommendations to address the challenges summarized above, namely:

- Review Volume 1 of the CSE and determine how well it works as a guide to writing resiliency requirements.
- Identify DoD/NSS [National Security Systems] policy references for identifying and understanding gaps, important for implementing the CSE.
- Participate in quarterly TEMs, as described in the roadmap below in Figure 6, in preparation for the 2017 Resiliency Invitational.

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

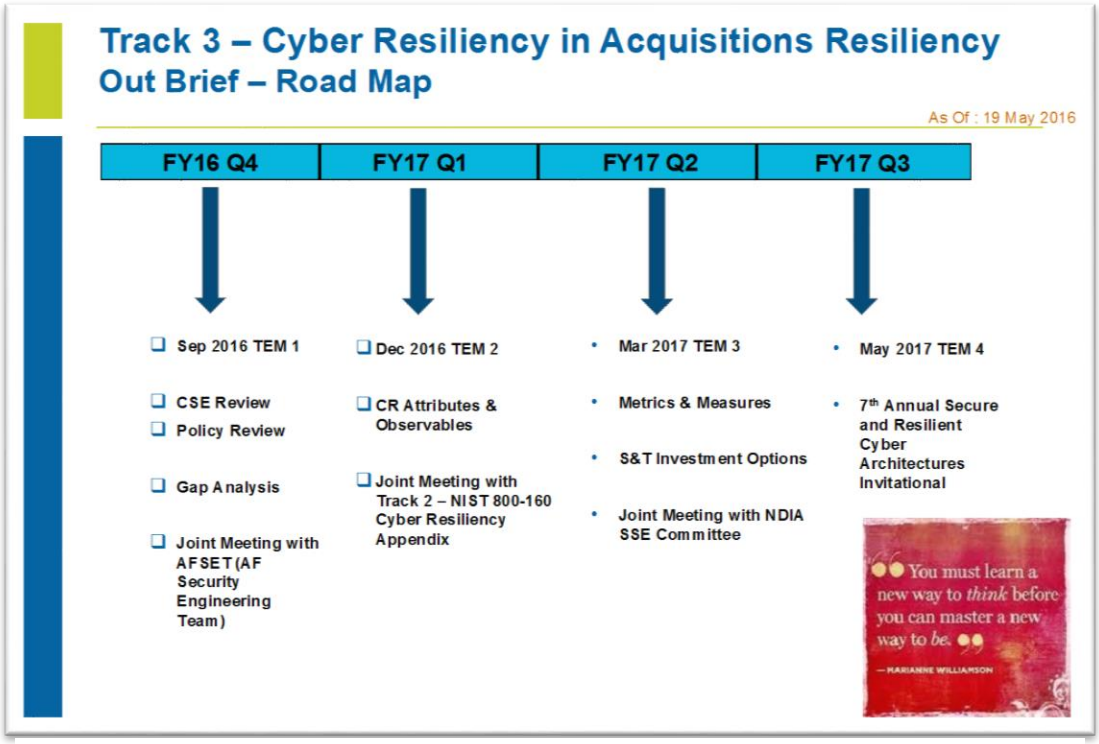

Figure 6. Road Map of Future Quarterly TEMs

This report was compiled from sessions at the MITRE Sixth Annual Secure and Resilient Cyber Architectures Invitational, 2016
Approved for Public Release; Distribution Unlimited, Case Number 17-0914.

References

[1] The MITRE Corporation (ed.), "Second Secure and Resilient Cyber Architectures Workshop: Final Report," 2012. [Online]. Available: https://www.mitre.org/cyberworkshop

[2] The MITRE Corporation (ed.), "Third Annual Secure and Resilient Cyber Architectures Workshop," December 2013. [Online]. Available: https://www.mitre.org/publications/technical-papers/third-annual-secure-and-resilient-cyber-architectures-workshop

[3] The MITRE Corporation (ed.), "Fourth Annual Secure and Resilient Cyber Architectures Invitational," 2015. [Online]. Available: https://www.mitre.org/cyberworkshop

[4] The MITRE Corporation (ed.), "Fifth Annual Secure and Resilient Cyber Architectures Invitational," 2016. [Online]. Available: https://www.mitre.org/cyberworkshop

[5] K. Heckman, F. Stech, B. Schmoker, and R. Thomas, "Denial and Deception in Cyber Defense," 2015. [Online]. Available: https://www.researchgate.net/publication/275270540_Denial_and_Deception_in_Cyber_Defense

[6] E. Hutchins, M. Cloppert, and R. Amin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Cyber Kill Chains," 2011. [Online]. Available: http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf.

[7] R. Lenz, "Cyber Security Maturity Model," 2011. [Online]. Available: http://www.dintel.org/Documentos/2011/Foros/ses2Mcafee/lentz.pdf

[8] The MITRE Corporation (ed.), "Industry Perspective on Cyber Resiliency," 2015. [Online]. Available: http://www2.mitre.org/public/industry-perspective/

Approved for Public Release; Distribution Unlimited, Case Number 17-0914.