Fifth Annual Secure and Resilient Cyber Architectures Invitational

Overview

The Annual Secure and Resilient Cyber Architectures Invitational (previously referred to as a workshop) accelerates recognition and adoption of cyber resilience by bringing together the cyber resiliency community for collective work on topics of common concern. 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, focused 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, focused on applying cyber resilience to space and critical infrastructure, designing a cyber resilience challenge, and identifying roles played by cyber resilience throughout the systems engineering life cycle [3].

The Fifth Annual Secure and Resilient Cyber Architectures Invitational, held in May 2015, focused on taking stock of the state of cyber resiliency – the lessons learned and the remaining challenges to overcome. The meeting attracted more than 150 attendees from government, industry, and academia. In addition, 11 vendor booths displayed cyber resiliency offerings. The first day of the Invitational featured keynotes and briefings by government, industry, and academic leaders to provide a common frame of reference and stimulate discussion of new ideas. The afternoon included a panel discussion, working group sessions, and Birds-of-a-Feather sessions. On the second day, the working group sessions continued and the day concluded with readouts and a way forward from each of the working groups.

This report presents a summary of the keynote talks and captures the goals, discussions, challenges, and recommendations of each working group session. Other materials from the Invitational and briefings can be found at http://www.mitre.org/cyberworkshop. We welcome comments from readers through the contact email address: secureandresilient@mitre.org.

The Cyber Resiliency Invitational Committee
April 2016

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

Acknowledgments

The MITRE Cyber Resiliency Invitational Committee members would like to acknowledge Nick Multari and Chris Oehmen from Pacific Northwest National Laboratory for their contributions to the Cyber Resiliency in the Cyber Physical Systems section of this paper.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

Table of Contents

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

List of Figures

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

1. Introduction

The Fifth Annual Secure and Resilient Cyber Architectures Invitational brought the cyber resiliency community together to accelerate recognition and adoption of cyber resilience. This fifth Invitational sought to produce community consensus around the theme of *Cyber Resilience: Looking Backward (What Has Worked? What Has Not?), Looking Forward (What New Challenges Must Be Faced?)*. The bulk of the first day's presentations set the theme for the event through the four keynote presentations. The keynote speakers included representatives from the National Institute of Standards and Technology (NIST), Navy, Indiana University, and Bit9 + Carbon Black. Section 2 summarizes these talks.

The afternoon of Day 1 and all of Day 2 consisted of facilitated working groups in four tracks. Sections 3–6 provide details about the following four tracks:

- Track 1: Analysis Through a Resilience Lens
- Track 2: Addressing Cyber Resiliency in a Traditional Cyber Security Focused World
- Track 3: Cyber Resiliency in the Cyber Physical Systems
- Track 4: Cyber Resilience Tabletop

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

2. Keynotes

The first day started with four keynote speeches. Dr. Ron Ross, Fellow at the National Institute of Standards and Technology (NIST) led the discussions with a presentation about an approach for building more trustworthy and resilient systems. Next, Juliana Vida, Deputy Director of the Department of Navy Deputy Chief Information Officer (Navy) (DDCIO(N)) Directorate within OPNAV N2N6, spoke about the Navy's Task Force Cyber Awakening. Then Craig Jackson, Senior Policy Analyst at Indiana University's Center for Applied Cybersecurity Research, discussed resiliency from a newcomer's perspective. Harry Sverdlove, Bit9 + Carbon Black's chief technology officer, delivered the last keynote, which focused on building resilient security.

2.1 Rethinking Security from the Inside Out: An Engineering and Life Cycle-Based Approach for Building More Trustworthy and Resilient Systems: Dr. Ron Ross

Dr. Ron Ross stated that our desire to implement advanced technology is rapidly exceeding our ability to protect the technology. Today we approach security bottom up instead of top down, outside in versus inside out, and tactically rather than strategically. Dr. Ross introduced the NIST Special Publication (SP) 800-160 *Systems Security Engineering* document and walked through an example that demonstrated the integration of security best practices and the value of the NIST SP 800-160. His final comments emphasized the need to be proactive in protecting organizational assets from cyber threats and he stated that "Security should be a by-product of good design and development practices – integrated throughout the organization."

2.2 U.S. Navy's Task Force Cyber Awakening 101: Ms. Juliana Vida, Deputy Director, DDCIO (N) Division / OPNAV N2N6BC

Juliana Vida described the U.S. Navy's Task Force Cyber Awakening. Due to the sensitivity of the presentation the information will not be shared on MITRE's Secure and Resilient Cyber Architectures Invitational Website or in this publicly available report. For more information, see [4].

2.3 Perceptions, and Barriers to Resilience: A Newcomer's Perspective: Craig Jackson

Craig Jackson examined the barriers to putting resilience into practice: *What definition is correct? Is it too advanced or too expensive? Is it really going to help?* He observed that even when people have a clear understanding of the ideas behind resilience they still have many other questions: *Is this strictly for advanced programs and professionals? Are there pre-requisites like having a mature program as a basis? What is the investment?* He suggested that the cyber resilience community provide newcomers with some direction, for example. by clarifying the audience for the resources, research, and products as well as providing roadmaps on how to get started. He also suggested that academia and research can make a huge contribution by

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

developing a body of empirical work around resilience in order to help the resiliency community gain a better understanding of resilience and prove the impact of the solutions we offer.

2.4 Building Resilient Security for the Age of Continuous Attacks: Harry Sverdlove

Harry Sverdlove noted that the technology landscape is continuously changing and vulnerabilities continue to rise each year. Threats evolve constantly, and a continuous stream of data breaches occurs every year. Cyber resiliency must involve continuous processes and security must also evolve continuously. In the security lifecycle of prevent, detect, and respond, most organizations spend the bulk of their budget on prevent and remain weak in predictive security. For detection, organizations still rely largely on point-in-time scanning and emphasize actionable threat intelligence. However, responses are not at all continuous and are expensive, reactive, and disruptive. Organizations should prioritize data collection in order to accelerate response time. At the end of his presentation Sverdlove mentioned the needs for constant detection and response, not just constant prevention.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

3. Analysis Through a Resilience Lens

Track Chair: James Cebula, SEI CERT
Track Scribe: Deborah Bodeau, The MITRE Corporation

3.1 Goal

This track provided a venue for participants to share experiences and lessons learned from performing assessments or analyses – of systems, programs, missions, complex and adaptive systems (typically systems of systems), and organizations – that focused on cyber resiliency or included cyber resiliency as one aspect.

3.2 Discussion / Observations

Multiple organizations have found that resilience-informed or resilience-focused analyses provide value, enabling them to improve the resilience of their operations and of the systems on which they depend. When an analysis is resilience-informed, the resilience perspective can provide an alternative view (whether of an engineered system or of a complex adaptive system). One type of resilience-focused assessment relies on an underlying model, such as the Computer Emergency Response Team's (CERT's) Resilience Management Model™ (RMM, [5]) or a maturity level model. Such model-based assessments provide the benefits of structure, consistency, and repeatability, but often lack a "hook" to grab the attention of senior leadership in an organization. Scenario-based risk assessments, when performed in conjunction with such model-based assessments, can illuminate the value of model-based recommendations.

Many different types of analyses can be informed by a cyber resilience perspective, including those focused on security, operational resilience, safety, survivability, or risk. Each type of analysis provides its own way to look at a system, relying on specific types of information and (implicitly or explicitly) on a system modeling approach. A cyber resilience perspective can be applied as an adjunct to any type of analysis. However, to include that perspective, additional types of information may be needed, and the relationship between the chosen system modeling approach and cyber resilience-related models (e.g., the Cyber Resiliency Engineering Framework or CREF [6]; a cyber kill chain or cyber attack lifecycle model) must be defined. In addition, any type of analysis has its own structural questions, such as how often should it be performed, or how much lead time is needed for the decisions to be informed by the analysis. The answers to those questions can strongly affect the level of detail possible in the resilience aspect of the analysis.

The scope of a resilience-informed or resilience-focused analysis depends in part on the level of abstraction at which the system (or system of systems) to be considered is defined. For example, there is a distinction between an information and communications technology (ICT) system that fuses and analyzes data from a variety of supervisory control and data acquisition (SCADA) systems, and a critical infrastructure system (more accurately, a system of systems) e.g., for water management, that depends on the output of that ICT system to determine how much water to deliver where. The resilience of the ICT system is intended to ensure the resilience of the

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

water management mission, while the critical infrastructure system as a whole must be resilient in order to deliver good to its end users. In terms of the multi-tiered approach to risk management defined in NIST SP 800-39 [7] (see Figure 1), the infrastructure system is identified with Tier 1 (the organization – or the water sector viewed as an organization), while the ICT system itself is a Tier 3 entity, and the combination of the ICT and the SCADA systems that inform it constitutes a Tier 2 entity (identified with a mission or business function).

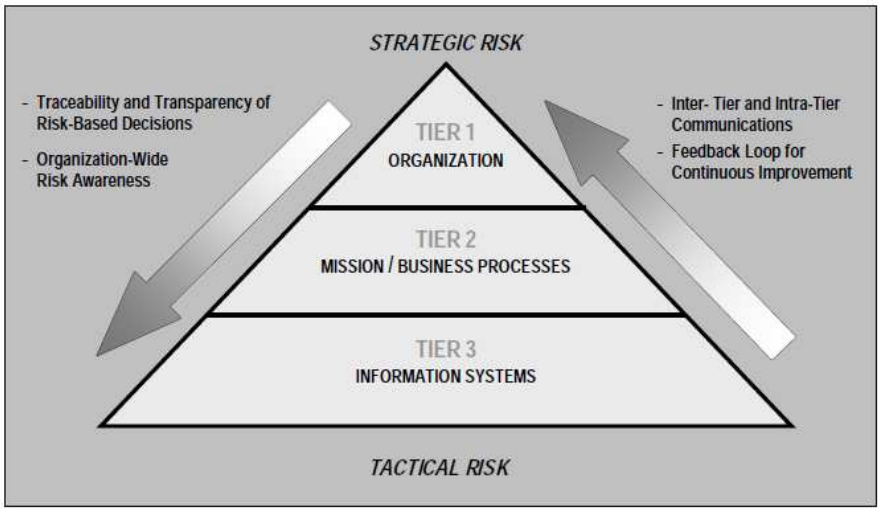

Figure 1. Multi-Tiered Approach to Risk Management (Source: NIST SP 800-39)

At the highest level, an analysis could look at an ecology of socio-technical systems. When the scope includes more separately managed systems – for example, when complex and adaptive systems are analyzed – the analysis must consider how resilience of one element (at a given level of abstraction) affects the resilience both of the whole and of other elements at the same level of abstraction. In particular, analysts must pay attention to the potential for cascading effects.

The scope of an analysis must be limited, or the analysis will be intractable. However, limiting the scope does exclude information, and requires that the analyst make assumptions. Experience has shown that unless assumptions and definitions of what is out-of-scope are made explicit, someone will assume a broader scope and a more completely informed analysis. The scope can be limited in multiple dimensions, including the level of abstraction or scale (ranging from a sector, an organization, a mission, a shared service or function, to a system or component), time, physical location, and governance. But it is crucial not to limit the set of challenges that resilience must address. While analyses have often constrained the stresses or types of disruptions (e.g., by focusing on previously observed events) a continually changing cyberspace will always present unforeseen challenges.

A resilience-oriented analysis can either distract or help decision makers. The results of an analysis can be distracting when priorities cannot be established (e.g., when stakeholder values are in conflict) or when the analysis appears to create a new set of problems that divert resources from already-recognized problems. The results of an analysis can be helpful when based on a prioritization scheme that is consistent with the organization's overall risk management strategy

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

(e.g., life-essential, safety-essential, mission-supportive). When consistent stakeholder values exist, analysis can help with prioritization. (Note that "consistent" can include "competing" – for example, stakeholders can compete for resources within a common framework of organizational missions, regulatory requirements, and budgetary processes.) Performing an analysis can help stakeholders to establish, discuss, and defend trade-offs, for example between efficiency and resilience, at all three tiers (e.g., policy trade-offs at Tier 1, mission trade-offs at Tier 2, and engineering or acquisition trade-offs at Tier 3).

Analyses can draw relevant experiences from established disciplines such as safety and security with which cyber resiliency conceptually overlaps and may have synergy. A key lesson from both those disciplines is that the desired property must be baked in rather than bolted on. Safety-critical systems also teach the lesson that it is vital to make the value proposition clear; for example, in addition to reduced risk, the benefits of improved safety (or resilience) can include lowered maintenance costs due to reduced unnecessary complexity. A key lesson from security is that a compliance orientation does not reduce risk cost-effectively; risk-based analyses that show how security contributes to mission assurance provide meaningful motivation for security measures.

Analyses can engage stakeholders in multiple ways. Experience has shown that bringing together a small community of highly engaged stakeholders to participate in an analysis is more effective than trying to engage a large group. Some stakeholders have a deep understanding of the problems they face; as participants in an analysis, they have much to offer. Ideally, an analysis would include participants with varied backgrounds or roles (e.g., risk manager, operator, end user, cyber defender) to represent the full spectrum of stakeholders. However, many stakeholders are already overwhelmed, do not want to be assigned another job, and even when motivated don't have the time or energy to participate effectively. Thus, in identifying participants to engage in a resilience-informed analysis, organizations must consider the full range of relationships. There is no one right relationship. However, a relationship management strategy (e.g., Responsible, Accountable, Supporting, Informed, Omitted) can be applied.

Cyber resiliency analysis involves sense-making under ambiguity. Two frameworks related to sense-making and risk governance can be used to help stakeholders think about cyber resiliency. The Cynefin framework [8] (see Figure 2) identifies problem domains as obvious, complicated, complex, or chaotic, with a cliff between the chaotic and obvious domains (organizations that believe cyber is simple will keep falling off the cliff). Note that the Cynefin framework has been applied to agile system development, and some thought has been given to its relationship to complex adaptive systems [9] and resilience [10].

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

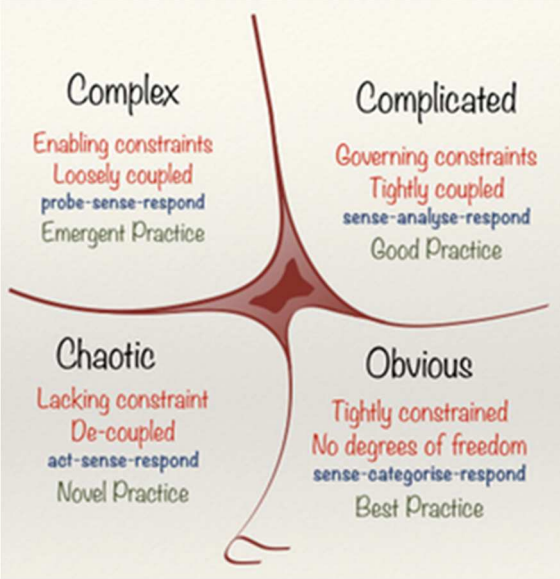

Figure 2. Cynefin Framework (Source: Wikipedia)

The International Risk Governance Council (IRGC) has defined a framework [11] for determining the range of participants who should be engaged in thinking about and making decisions related to risk management, based on the dominant characteristic of the problem domain: simplicity, complexity, uncertainty, or ambiguity (see Figure 3). Despite the slightly different interpretation of "complexity," these classes of problem domains are basically the same as those in the Cynefin framework.

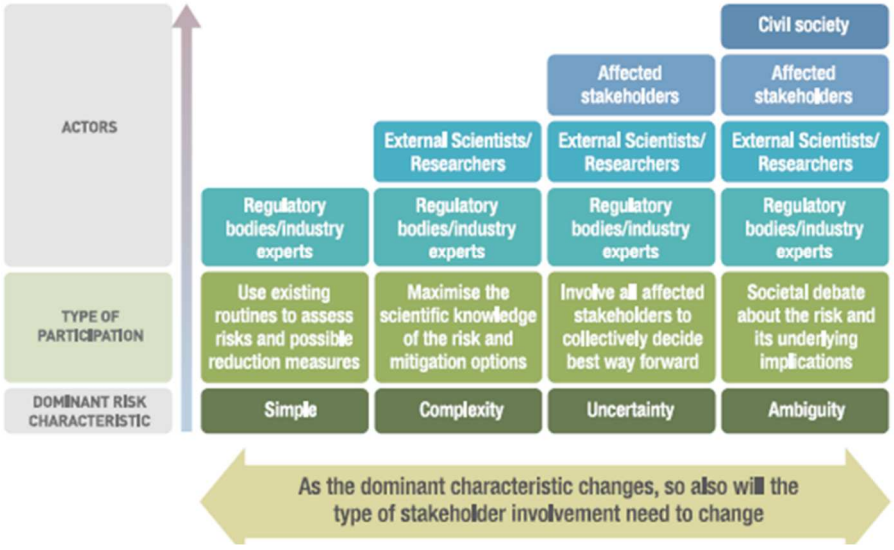

Figure 3. IRGC Framework

Achieving cyber resiliency – and, more broadly, operational resilience – involves leveraging relationships across the physical, information, cognitive, and social domains. The Net-Centric

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

Operations Conceptual Framework [12] defines such relationships. Security and cyber resiliency objectives can be mapped to the goals of protecting services, exchanges, and information. The framework provides a way of looking at quality and agility, where quality applies to networking, interactions, and information, and agility includes command and control (C2) and force agility, with an emphasis on synchronization.

3.3 Challenges

Identified challenges relate to terminology, stakeholders, and culture.

The problem of *terminology* has been raised in prior Invitational workshops. The term "resilience" is often used without definition; while various policy documents have defined it, those documents may have only minimal relevance to any given organization. The term "cyber resilience" is not defined in any policy or standard, and is often used without definition. (Note that in the Invitational workshops, the advanced persistent threat (APT) or advanced cyber threat has always been a consideration, while the specific definition has evolved over time.) The terms "organizational resilience" and "operational resilience" are well defined in the context of the RMM, but are often used outside that context. When organizations consider performing a resilience-informed or resilience-focused analysis, the planning must include communication and education about resilience and defining terminology in ways that are meaningful to the stakeholders who will be involved.

Organizations need to engage a variety of *stakeholders* in a resilience-informed analysis. (The specific set of stakeholders depends on the scope and purpose of the analysis, and on the governance structure in which the analysis is intended to inform decisions.) Stakeholder input is crucial, both to shape the analysis and to provide the information needed to support the analysis; their buy-in is also crucial, since they will be involved in implementing recommendations. However, stakeholders vary widely in their levels of understanding, both of the systems for which they are responsible (or on which they rely) and of potential adversity, particularly of prolonged campaigns by the APT. To the extent that an organization lacks a culture of shared concern for and responsibility for operational and cyber resilience, stakeholders cannot engage effectively.

Culture affects multiple stakeholders, in different ways. In system acquisition or evolution, legacy performers (contractors, developers, program office staff) are invested in existing architectures and processes. For critical infrastructure sectors, regulators and legislative actors are invested in existing ways of understanding the problem domain and solution space (and thus may be biased against considering the APT), and in existing regulations, standards, and evaluation processes. In many organizations, decision making is shaped by a just-in-time or near-term bias. However, organizations must exercise some degree of patience and restraint to make progress in resilience (such as improved security or safety). It is possible to perform analyses focused on "low-hanging fruit" or "quick wins," but the recommendations must be crafted to avoid distracting from the long-term value; instead, recommended near-term actions should demonstrate enough value to provide traction for long-term changes.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

3.4 Recommendations/Way Forward

Overall, the presentations and discussions in this track noted that:

- Resilience in general, and cyber resiliency in particular, can be the focus of or one aspect of an analysis.
    - Resilience-focused analyses produce recommendations specifically intended to improve resilience.
    - Resilience-informed analyses produce recommendations that are better grounded in real-world concerns than analyses that do not consider resilience.
    - Resilience-oriented (resilience-informed or resilience-focused) analyses can be performed for different scopes, timeframes, and levels of abstraction.
- Stakeholder engagement is crucial to the effectiveness of resilience-oriented analyses.
    - Sense-making and risk governance frameworks provide a rationale for representing a wide range of stakeholders in resilience-oriented analyses.
    - Common models or frameworks help stakeholders make sense of the problem and solution domains together.

4. Addressing Cyber Resiliency in a Conventional Cyber Security Focused World

Track Chair: Tom Llanso, JHU/APL
Track sub-chair: Richard Graubart, The MITRE Corporation

4.1 Goal

Cyber resiliency is an area of growing importance as organizations recognize that it is not realistic to assume that conventional cyber security measures will always be successful in keeping adversaries out, or providing organizations the ability to quickly detect and eradicate any footholds by the adversary. Cyber resiliency cannot succeed without a good foundation of conventional cyber security. The question becomes what is the appropriate balance between conventional cyber security and cyber resiliency, and how does one recognize and reconcile the trade-offs between the two.

4.2 Discussion / Observations

While it is important not to turn this into definitional battle, it is also important that the concepts underlying the terms be understood. Cyber resiliency (also referred to as cyber resilience) is *the ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on cyber resources*.[1] Cyber resiliency differs from conventional cyber security in two ways:

- First, conventional cyber security[2] focuses on achieving the security objectives of confidentiality, integrity, and availability to acceptable levels, by using a combination of perimeter protections and internal controls. Cyber resiliency focuses on the cyber aspects of mission assurance by maximizing the ability of the cyber resources on which the missions depend to accomplish necessary tasks, even in the face of adversity – and thus on the mission-related objectives of anticipate, withstand, recover, and evolve.
- Second, the threat model for conventional cyber security assumes that the adversary can be kept out of a system or can be quickly detected and removed from that system. In contrast, cyber resiliency is based on the recognition that such an assumption is unrealistic. In light of supply chain contamination, long adversary dwell times, and loss of control over parts of the set of cyber resources on which an organization depends, the

[1] Cyber resources are "Separately manageable resources in cyberspace, including information in electronic form, as well as information systems, systems-of-systems, network infrastructures, shared services, and devices." - derived from NIST SP 800-39 [7].
[2] Conventional security is the primary focus of FIPS 199 [15] and the baselines defined in NIST SP 800-53R4 [16] and CNSSI 1253 [18]. The series of publications by the Joint Transformation Initiative (JTI) – including NIST SP 800-39 [7], NIST SP 800-53R4 [16], and NIST SP 800-30R1 [17] – include consideration of advanced cyber threats [19] and cyber resiliency, but organizations using those publications can restrict themselves to conventional threats. Similarly, the NIST Cybersecurity Framework [13] accommodates consideration of advanced threats at Implementation Tier 4, but does not provide guidance on whether and when an organization should seek to be at Implementation Tier 4.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

goals of cyber resilience move beyond the "Identify, Protect, Detect, Respond, Recover" model [13], although they complement that model.

Cyber resilience works in conjunction with cyber security. Most cyber resilience measures assume, leverage, or enhance a variety of cyber security measures. Cyber security and cyber resiliency measures are most effective when applied together in a balanced way. The cyber resiliency perspective reflects that modern systems are large and complex entities, and as such systems, operational environments, and supply chains will always have flaws and weaknesses that adversaries can exploit. Given resource limitations, it is not practical to try to maximize mitigations intended to enhance both cyber resiliency and conventional cyber security. Therefore, achieving effective protection of cyber components often requires a mixture of conventional cyber security and cyber resiliency investments. Placing too much emphasis on conventional cyber security measures designed to keep the adversary out may leave a system with insufficient cyber resiliency mitigations, thus making it unable to effectively counter and respond once an adversary has achieved a foothold. However, placing too much emphasis on cyber resiliency investments could make the security boundaries overly porous; this could result in a larger attack surface, enable adversaries to achieve effects immediately as well as to obtain and exploit a foothold, and cause cyber resiliency measures to be overwhelmed by an increased number of low-level attacks.

It is also important to recognize that the term employed in this section is "conventional cyber security." The modifier "conventional" is critical, as it reflects how cyber security is practiced today. Over time cyber resiliency should be included in a broader construct of cyber security as reflected in Figure 4.

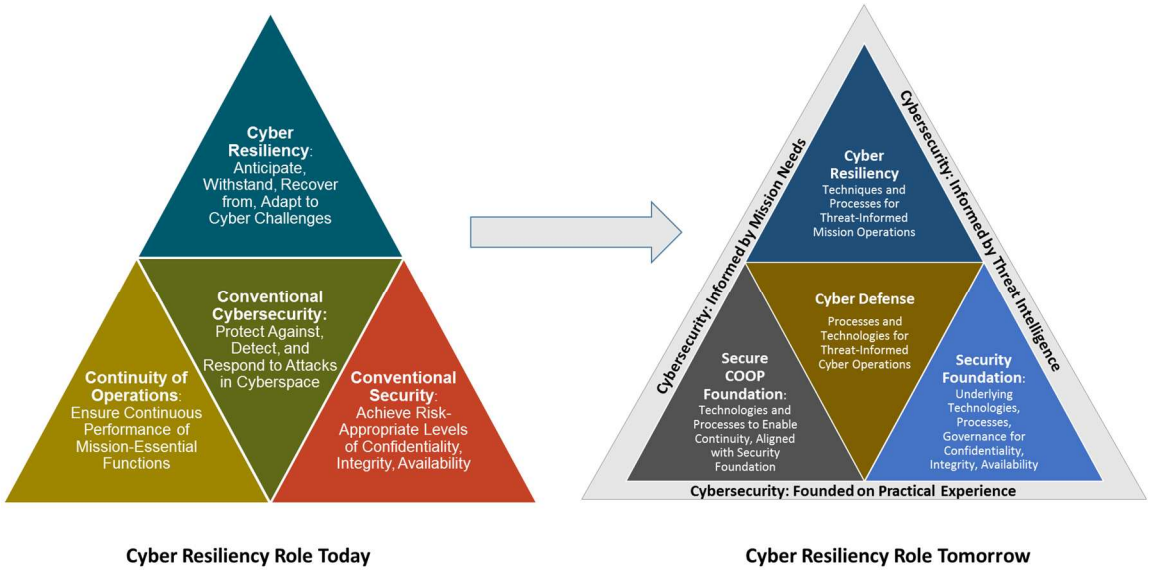

Figure 4. Changing Relationship of Cyber Resiliency and Cyber Security

A key point in considering cyber resiliency and conventional cyber security is that both should be viewed through the mission lens. Both represent means to an end, not ends in themselves.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

Cyber resiliency supports mission assurance and mission continuity by focusing on the cyber systems and components on which a mission depends. It aims to maximize an organization's ability to complete critical mission functions despite an adversary presence in the organization infrastructure (threatening mission-essential cyber components). In effect, cyber resiliency enables an organization to "fight through" despite an adversary's having achieved a persistent foothold in the organization's infrastructure.

Similarly, cyber resiliency and conventional cyber security are both closely linked to systems security engineering. When correctly implemented, systems security engineering will help implement both conventional cyber security and cyber resiliency. However, cyber resiliency often focuses on capabilities and emphasizes goals different from those *conventionally* emphasized in systems security engineering. For example, the cyber resiliency goals of anticipate, withstand, recover, and adapt complement the security objectives of confidentiality, integrity, and availability that are conventional goals of systems security engineering.

4.3 Challenges

The working group noted that trying achieve the appropriate balance of cyber security and cyber resiliency presents multiple challenges.

Education and Culture

One of the first challenges concerns education and culture. Unfortunately, many individuals fail to recognize that the conventional means of securing systems (e.g., firewalls, strong identification and authentication, keeping security software patched and up to date) are not sufficient to secure systems against the advanced cyber threat.

Still another, albeit related challenge, is that many believe that they can ensure cyber resiliency by simply following the NIST or Committee on National Security Systems Instruction (CNSSI) baselines, but these baselines only partially address cyber resilience. Of the approximately 150 cyber resiliency controls in NIST SP 800-53 [14], only about 60 of them are included in the CNSSI 1253 High Confidentiality, High Integrity, and High Availability (HHH) baseline (and a lower number and percentage in the other baselines). In other words, only approximately 13 percent of the HHH baseline's 460+ controls focus on resiliency. Moreover, the selected HHH resiliency controls focus largely on enhancing aspects of conventional cyber security (e.g., detecting the adversary, analyzing its actions, and recovering from those actions); those are important, but do not cover all aspects of cyber resiliency. NIST SP 800-53 contains cyber resiliency controls in, not selected in any baseline that focus primarily on disrupting the attack surface. This impedes adversaries' ability to get correct and timely information about defensive capabilities and causes them to make incorrect assumptions about the system or its defensive capabilities, waste resources, or prematurely disclose malware to cyber defenders.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

Compliance Mindset

A major issue (for both cyber resiliency and cyber security) is the mistaken belief that the baselines (NIST or CNSSI) are mandatory. In reality, the baselines are merely starting points;[3] organizations should diverge as appropriate (which includes adding or removing controls from baselines). An organization might choose to diverge from a baseline and replace it with a different set of controls (in some cases more, in other cases fewer). This mistaken view is relevant to cyber resiliency, because organizations with this mistaken belief might therefore think that achieving a greater degree of cyber resiliency requires that they address the controls in the baseline *and* additional cyber resiliency controls. In actuality, consistent with the risk management framework (RMF) and trade space analysis inherent in system security engineering, and because resources are finite, an organization might choose to withdraw resources from some more conventional cyber security mitigations (e.g., security controls focused on hardening the boundaries), and place greater emphasis on those mitigations (controls) employed in countering an adversary that has achieved a persistent foothold within the organization's systems.

Lack of Data and Measurements

A challenge specifically in regard to cyber resiliency centers on data and measurement. Cyber resiliency is a relatively new area and only a limited number of organizations have adopted it; as a result, we have an insufficient body of evidence to judge resiliency. Thus, while we have reasonable expectations regarding the types of effects a given resiliency technique[4] might have on an adversary, and which techniques might be better suited to one environment than another, we lack concrete metrics regarding the efficacy of any cyber resiliency techniques (and corresponding cyber resiliency controls) on an adversary.

The problematic nature of the current lack of good cyber resiliency metrics is compounded by the tight linkage between cyber resiliency and mission assurance (continuity). Thus the true measure of cyber resiliency effectiveness comes from judging the ability of missions to continue to operate despite ongoing adversary presence.[5] In effect, it requires that organizations incorporate cyber resiliency into systems and use it to address the advanced cyber threat in order to obtain a true measure of the efficacy of mitigations that support cyber resiliency.

This lack of data impacts trade space analysis. Cyber resiliency is not free and organizations have a finite amount of resources available that they can apply to securing their systems,

---

[3] "The use of the term baseline is intentional. The security controls and control enhancements in the baselines are a point of departure from which controls/enhancements may be removed, added, or specialized based on the tailoring guidance…" Section 3.1 NIST 800-53 R4 [16].

[4] Cyber resiliency techniques characterize approaches to achieving one or more cyber resiliency objectives that can be applied to the architecture or design of mission/business functions and the cyber resources that support them. Each technique refers to a set of related approaches and technologies [20] [6].

[5] This topic was discussed in more detail in the Metrics and Assessment track of the 2012 Secure and Resilient Cyber Architectures Invitational [1]. Metrics related to cyber courses of action were discussed in the Active and Adaptive Response track of the 2013 Invitational [2]. The need for cyber resiliency metrics was also affirmed by the Systems Engineering Lifecycle track of the 2014 Invitational [3].

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

including to cyber resiliency. However, without concrete data, systems engineers cannot definitively identify the optimum mitigations that support cyber resiliency, and the appropriate trade-offs between applying resources to cyber resiliency as opposed to conventional cyber security.

Malware Detection Example

> We know that signature-based anti-virus software, detonation chambers, and honeynets are all capable of detecting adversary malware. Anti-virus software is of course common practice (classic cyber security), and the concept of employing anti-virus software from multiple vendors (e.g., one at firewall, another at email gateway, still another on laptops) is gaining traction. The use of detonation chambers and honeynets is much less common, and when employed is used to detect malware that has breached an organization's boundary (example of cyber resiliency). But we lack concrete metrics regarding which means of detecting malware are more effective. Similarly, we lack information on the trade-offs between these measures; determining the optimum number of different anti-virus software suites, establishing whether money is more effectively spent on anti-virus software than on detonation chambers/honeynets, etc.

4.4 Recommendations/Way Forward

The working group had a variety of suggestions for moving forward. Many of them focused primarily on cyber resiliency *per se* rather than its relationship to cyber security. That reflected the group's recognition that the area of cyber resiliency is still relatively new.

Training and Guidance

Decision makers need training to better understand cyber resiliency and its relationship to cyber security. The training could take the form of classes, briefings, or documents. Whatever the form, it should be appropriate for the audience – guidance for mission owners will be different than that for program managers (PMs), than that for warfighters, etc.

It is important to realize that it is not sufficient to simply direct decision makers to documents such as NIST SP 800-53, 800-39, etc. Those documents, while excellent, are geared toward individuals who have significant involvement in and understanding of cyber security. High-level documents, such as slip sheets, are needed to clearly, cogently, and concisely convey key cyber security and cyber resiliency information to decision makers. In addition, "how to" guides are needed that would help organizations 'get started' on advancing cyber resiliency. Since such guides would be designed for organizations that might not be well versed in cyber resiliency the content should be:

- Broken down by audience types
- Broken down by mission / system types
    - E.g., enterprise IT environments, cyber physical / tactical systems

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

Activities to Facilitate a Change from Compliance Mindset

It is essential to shift the mindset of individuals away from compliance and toward the concept of risk management. This change is broader than just cyber resiliency: it is consistent with the RMF and some of the key underpinnings of system security engineering. Fundamental to this shift is emphasizing both in doctrine and practice that published baselines are simply starting points, not mandates.

To support risk management trade-offs, stakeholders need guidance to help them make reasonable, well-funded, and ideally repeatable risk management decisions. The group recommended that a repository of cyber resilience information should be developed and kept current. Such a repository should include:

- Relevant cyber resiliency frameworks
- An updated listing of controls in NIST SP 800-53 that support cyber resiliency;
- A resiliency 101 document roadmap, walking readers through various cyber resiliency material
- Pointers to cyber resiliency conference proceedings, including this cyber resiliency workshop
- Use cases or worked examples of applying cyber resiliency;
- A short, easy-to-understand list of cyber resiliency best practices
- A cyber resiliency FAQ.

Tools

Part of the challenge of performing cyber resiliency-related assessments and decisions stems from the lack of tools to support such decisions. Going through several dozen controls is a labor intensive activity. Among the type of tools needed are:

- Automation that would allow organizations to enter the characteristics of their environment (e.g., tactical), assumptions associated with a controls (e.g., assumes persistent data/services), overlays/baselines in which controls appear, and whether a control supports cyber resiliency. The ability to query such tools would allow a system security engineer to relatively quickly identify the appropriate resiliency-related mitigations for their system. Note that this tool would not replace trained system security engineers, but merely supply a means to allow them to do their job faster.
- Automation of the OODA loop, with the adversary included in the loop, and identification of which aspects require human actions and which do not
- Continuous exercising of production environments via automated tools (e.g., Simian army), which would allow for the testing of the efficacy of various cyber resiliency mitigations.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

Measurements and Data

A better understanding of cyber resiliency-focused mitigations, the efficacy of those mitigations, and their relationship to cyber security mitigations requires well-founded data and associated measurements regarding the effectiveness of cyber resiliency measures. Achieving this requires,

- Identification and application of fundamental cyber resiliency measurements
- Linkage between cyber resiliency measurements and mission assurance measurements
- Collection of experience of applying cyber resiliency in operational systems. Unfortunately, given the slow procurement process, considerable time may elapse till sufficient cyber resiliency information becomes available from operational systems. Therefore, it is likely that in the immediate future this information will only be available via modeling and simulation.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

5. Cyber Resiliency in the Cyber Physical Systems

Track Chair: Nick Multari, PNNL
Track Sub-Chair: Chris Oehmen, PNNL

5.1 Goal

This track extended prior work done during the Resilient Critical Infrastructures track at the MITRE Secure and Resilient Cyber Architectures Invitational in 2014. In the 2014 track, the participants developed a consensus working definition of resilience as "the extent to which a complex cyber system can continue to deliver critical services in the face of impediments." They then developed four profiles (communications, finance, energy, transportation) for representative critical infrastructure companies and explored the barriers to achieving resilience in these exemplar companies. The 2015 Cyber Physical track used these four exemplar companies and the working definition of resilience to take this process one step further, exploring issues around resilience in a subset of the companies that have significant reliance on cyber physical systems. The goal is for these to serve as a model of how resilience might be realized for cyber physical systems in other critical infrastructure domains.

5.2 Discussion / Observations

The members in this track first sought to develop a working definition of "cyber physical systems." Though many people have an intuitive sense of what the term means, for the purposes of analyzing the cyber physical systems and dependencies within critical infrastructure domains it is important to distinguish between what the broad definition of cyber physical systems includes and what it does not. To that end the group explored four different types of definitions.

1. **Attack-focused definitions:**
    - Cyber attack that creates physical impact.
    - Physical and cyber attack that together can cause more damage than either one on its own.
2. **System-focused definitions:**
    - Physical system (car, satellite, etc.) enhanced by a cyber system that also introduces an attack surface.
    - System that has both cyber and physical components. The emphasis here should be on systems where some physical component carries out a mission, and is supported by a cyber infrastructure. As an example, an industrial control system (ICS) has to turn or move something, but is connected to a cyber component, and therefore can be attacked through that vector.
3. **Process-related definitions**
    - Some kind of stimulus/response that takes place between the physical and cyber environment, where the response either controls or is controlled by cyber (feedback loop).

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

4. **Open source definitions**
    - NIST: "computational processes that interact with physical processes that are tightly interconnected and coordinate to work effectively, often operating with humans in the loop." Some examples are:
        - Smart manufacturing and production
        - Transportation and mobility
        - Energy
        - Civil infrastructure
        - Health care
        - Buildings and infrastructure
        - Defense and response
    - DHS: "coupling of physical and information/processing, specifically when one can't decouple these and still get the intended outcome." This view does not focus on fully automated processes but instead emphasizes the human in the loop. Under this view, there is an elevated possibility for unexpected emergent behaviors.
    - Wikipedia: "system of collaborating computational elements controlling physical entities." The group noted that this definition is highly influenced by an embedded systems perspective.

The group decided to adopt as a working definition for cyber physical systems "tightly coupled systems involving both cyber and physical components, possibly with a human in the loop." Physical effects are the output of the system, and this physical output is an essential feature of the system's mission. In this view, communications is an enabler, finance is a consumer, but neither is a cyber physical system by this definition. *The result was to reduce the four critical infrastructure exemplar companies to two and focus on transportation and energy for the remainder of the track.*

Cyber Physical Systems & External and Internal Dependencies

Transportation: The exemplar transportation company explored in this track includes trucks, planes, people, systems, and drones. The main function is to deliver packages efficiently and affordably. This is its core business, from which it obtains competitive advantage, so it must be one of the key functions provided by the cyber physical systems. This requires logistics, schedules, maintenance, command and control, and situational awareness (e.g., where packages, people, and vehicles are; dynamic road and weather conditions; package handling systems). Every distribution node communicates with a cyber system since each hub or super hub communicates information about package state. Important external dependencies include communication systems, other delivery and pickup partners, airports, fuel, and utilities.

Key cyber physical systems in this domain are:

1. Distribution system (move/sort/track packages),
2. Delivery drones, and drivers with scanners and other edge points for the system.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

These cyber physical systems have critical dependencies on both guarantees of the reliability of the command and control network and preexisting systems of unique alphanumeric text identifiers for worldwide endpoints. These identifiers were identified as particularly critical because they serve as the bridge between machines and humans.

Key interdependencies between transportation and energy are fuel pipelines, infrastructure to support smart electric vehicles, and gas trucks/tankers. It is important to note that this doesn't deal with hazardous transportation, which would bring in new requirements for resilience as logistical constraints. For example, packages can sit on a tarmac for long time unless they have hazardous materials.

Energy: For the energy company, the group identified several internal cyber physical systems. These include:

1. ICS/SCADA components such as circuit breakers, synchrophasors, and automated control systems
2. Sensors for predictive analytics such as data going to a control center to inform human decision makers
3. Smart grid components with intelligent control systems making local decisions
4. Switches that control failover for backup systems
5. Control systems for nuclear power plant fuel rods that pull up or insert the rods based on power sensing
6. Systems that control exhaust
7. Control of wind turbines based on wind speed

Several external dependencies were also identified, including fuel delivery; water to drive dam turbines, such as a steam source for steam-electric plants, and ??? for nuclear plants; and communications systems required for synchronization of control.

Resilience Requirements of Cyber Physical Systems in the Context of each Mission

Transportation: The demand for on-time delivery often dominates the need for resilience in this critical infrastructure sector. Secondary missions include special handling and responding to customer queries, with the constant need to do all this at minimum cost. In the event of a cyber attack, it will be critical that no disruption to public space or safety results from failures in the transportation system. Particularly vulnerable aspects of the transportation industry in the context of cyber attack include routing information and controllers responsible for sorting, whose failure could result in many failures of the primary mission of on-time delivery. Enabling resilience in this sense would require detection (situational awareness) as early as possible when the system is not operating correctly. This might be done in the future by comparing some key external measurements to expectations. It might also involve a human in the loop.

Technically the company might obtain this sort of situational awareness through actively perturbing the system by injecting tracer packages or rogue packages. The measured impact of these packages would give some idea of system performance. It will be important to differentiate

between when certain information is needed and when it is not. Put another way, requirements for the expected behavior of the sensory system must be matched to each different mission—5-day packages do not have the same resilience needs as next-day air.

Economic factors also drive interesting resilience requirements. For example, knowing that resilience technologies may not be universally deployable at a single time point, the company should have a process for deploying resilience technologies and improvements in a staged order, with maximum increase in benefit at each stage. To optimize this process and other resilience processes, the company must recognize the need to protect against different classes of impediments in pure cyber systems or pure physical systems.

The group identified several social concerns with resilience in transportation systems, including (lack of) public acceptance of drones and driverless cars.

Energy: The key resilience needs for cyber physical systems in the energy sector are failover and redundancy to assure maximum availability of energy. Possible technical resilience strategies include redundancy of control systems, particularly combined with diversity of operating systems. Another good strategy might consist of having hot spares online so that the systems could revert to a good prior state in the event of problems. But some infrastructure. Such as power lines, cannot be duplicated and/or diversified to support a system-wide resilience strategy.

Some economic barriers to resilience relate to constraints imposed by what companies can afford and what customers would be willing to pay for. Cost associated with any change to power generation or delivery translates into additional end user cost.

Policy issues exist around moving to service providers. For instance, currently most of the emphasis is on availability, but no existing policy regulations direct utilities to address cyber concerns. If a cyber attack occurs, how should a service provider determine whether the producer, operator, or end-user is responsible or liable? The answer to this question will have large ramifications on who implements cyber resiliency controls and how that is done.

Some technical concerns center on cyber physical system resilience in the energy sector. For instance, if software is a service, the code driving cyber physical systems could be either owned or licensed, with large implications for the extent to which redundant or diverse resilience could even be implemented. Additionally, a lack of transparency might make some parts of the system opaque to inspection and control.

The main social issues center around the interplay between any additional regulation and the need to answer to policy makers and the population on issues such as raising rates. Redundancy may be good for resilience, but social pressures will resist building new plants. There is also social resistance to smart grid or other intrusive control from providers. Buy-in may have to be incentivized via cost savings.

A key social issue for both critical infrastructure domains is the problem of supply chain. Both industries must trust components without having control of their supply chain. In many cases

seemingly diverse supply chains actually rely on common components, sources, or technologies, so without knowledge of the whole supply chain it may be impossible to ensure redundancy that includes diversity.

5.3 Challenges

Technologies Needed for Resilience of Cyber Physical Systems and Barriers to Acceptance

Transportation: Key technology needs in the area of transportation include resilient hardware, databases, algorithms, and software. The company first needs to ensure that all data is correct. If addresses, routing, or other data is corrupted, the whole system does not work. Independent methods for verifying key data are highly desirable, including checks on the mapping between human-readable data such as addresses and bar codes that are machine readable. Once data is in the database, it must remain protected.

Regarding resilience in routing and command and control channels for transportation, it will be important to discover when a route sequence contains an error and how to re-route the item to ensure successful delivery. For example, if a package arrives at a particular stop later than expected, what logic should govern choice of the next step, for example to favor air instead of ground to compensate for the error? This problem is exacerbated when natural events such as earthquakes interfere with collocated paths that may otherwise constitute the redundant solution. In such a case, the detection of package delay is the first step, and ensuring correctness of future steps then becomes the goal of resilience.

The working group members shared a strong belief that new technology may not be necessary to enable resilience in cyber physical systems. In many cases resilience may be attained by connecting existing components using different principles and better understanding the attributes of those existing technologies. It will be important to develop algorithms to decide when a decision is "good enough" and "correct enough," because completely certain information may not always be attainable in a reasonable time. Combined modeling of cyber and physical components may be a critical enabler of this concept.

Several challenges complicate enabling resilience in cyber physical systems, largely driven by the sheer complexity of these systems. Not only are the components of these systems highly heterogeneous and distributed, but humans and other processes also introduce errors that may result in emergent behaviors. Additionally, much of the infrastructure may not be owned by the same entity and may be geographically disbursed. Users and operators will have to ensure that resilience technologies operating in failure modes in different geographical locations do not create legal problems. For example, if a plane lands on time in country X, but leaves late because of a problem in country Y, which affects an end user in country Z, who is liable when the three countries' laws are not consistent?

One primary concern for technology adoption is the loss of data privacy that may be required to ensure resilience. Often the actual connection between the distribution network and the physical

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

handing of a package to a customer is outsourced, creating a situation where possibly untrusted third-party employees have access to data, packages, or the package recipient.

Energy: Resilience technologies for the energy sector include operating system diversity for critical components, continual refresh of those operating systems by restoring them to a known good state, employment of more isolated components with out-of-band communications, and self-healing data through technologies such as ZFS—a file system that constantly double-checks file integrity and repairs detected errors.

The working group identified several barriers to acceptance of resilience that are common to multiple cyber physical systems, including skill level of the staff, size of the staff, cost of investing in hardware/software, maintenance and upgrade cost, and lack of specific resilience-related training for staff. The group also identified several challenges not common to all cyber physical systems and are especially important for the energy sector. The first centers on locations of distant components. Many of the cyber physical components in energy systems are in geographically isolated locations, so any resilience-related response may require resources that are disproportionately burdensome compared to more centralized cyber physical systems. Energy systems also have a strong regulatory mandate. Regulation of the power grid in the United States is highly localized and fragmented. By comparison, in the United Kingdom this function is instead centralized under a public utility commission. One possible solution in the United States would be to have representatives of the Department of Energy or another national-level entity accompany utilities at rate hearings to discuss cyber threat and technologies that would help normalize the cyber resilience response across these local regulatory boundaries.

Another problem arises from the lack of true valuation of electricity. This has implications for resilience in that the cost of electricity does not strictly correlate to its intrinsic value, so arguments that increased costs associated with elevated cyber resilience could be passed to the consumer may not hold.

One key social barrier to acceptance of resilience in cyber physical systems stems from the difficulty of obtaining consumer participation in a resilient response. For example, if recovery from an attack (or even a proactive response) requires deliberate loss of service that leads to a better overall degree of service, how would the company solicit consumer consent, and how would it communicate the need at all? Another social barrier arises because highly resilient systems make humans more reliant on them, thereby raising the severity of failures. This could actually lead to nonlinear effects in that humans are themselves part of the cyber physical systems, but capturing this for planning and system design may be very difficult.

One technical challenge identified for refreshing technologies was the need to protect “gold” copies. Any strategy that relies on the integrity and availability of a gold copy of data, applications, or operating systems simply creates a new attack surface. This strategy also presumes that enough state can be brought to bear to evolve a gold copy to a reasonably current reconstruction of a system, but it is not clear how that state would be preserved or trusted in the event of a system breach. The process of rolling back to a known good state also requires time and resources to recover. Failover and checkpoint restart may involve very lengthy processes and

therefore exceed the window of response time that is useful for energy applications. Additionally, some data used to maintain a gold copy comes from sensors. This relies on source of the data remaining uncorrupted, which imposes further constraints on the overall system. It is also not clear what the appropriate data retention policies would be for the short and long terms, including the feasibility of using offsite collection, backup, and storage. Moreover, energy companies have a strong dependence on external services such as the Global Positioning System (GPS) that may be beyond the bounds of the system and therefore outside the scope of control in a resilient response.

Companies confront a significant economic barrier in introducing cyber resilience to energy systems because many of those systems still run old operating systems such as WinXP or Win98. Rewriting legacy codes for newer systems or for alternative operating systems such as Linux to support diversity resilience applications is not always feasible.

5.4 Way Forward/Recommendations

Future Vision and the Key Steps in Realizing It

Transportation: Joint models of cyber and physical systems are the top priority for overcoming the barriers to resilience in cyber physical systems in transportation. The ability to model cyber physical systems across multiple critical infrastructures will prove highly important. As more and more consumers buy online, companies may have an opportunity to create a different business model that facilitates resilience; for example, by having smaller stores closer to end points as opposed to operating from large-scale centralized facilities. This may simplify the modeling required to understand the behavior of the overall system. The possibility of using drones for delivery may help avoid some entanglements of cyber physical systems, but it would certainly introduce new cyber attack surfaces.

Another emergent technology that may introduce new opportunities for transportation is additive manufacturing, or "3D printing." This may make it possible to produce or manufacture products even closer to their end points. However, some social issues would need to addressed for this business model to work. For example, it would be important to resolve territorial rights to determine who can fly where and how high. In addition, privacy and safety issues resolve around misrouting or errors in shipping that might be very different for drone delivery of additively manufactured products from distributed locations. This option also introduces new legal issues, for instance the need to determine what happens if an automated delivery system is hacked and causes property damage or personal injury.

One important cyber challenge to address centers on solving nonrepudiation issues to verify that delivery has occurred. One solution might be a beacon for each subscriber that they can be uniquely found, but this would introduce a new cyber attack surface. Companies would have to develop new resilient algorithms for routing infrastructure to ensure that the instructions to delivery drones are correct.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

One solution that would have significant impact in this area is active information. Currently, packages are identified using static codes with static information. Active monitoring of location by the packages themselves would produce situational awareness that could be used to provide redundancy in awareness and routing. Distributed awareness using local information used in this way would serve as independent validation.

Energy: New technologies that enable self-healing of cyber and physical components would greatly aid resilience for energy cyber physical systems. The key is that they would have to continue to operate even when critical infrastructure systems on which they depend are not available. This could be accomplished most directly using redundancy. For example, control signals that currently use wireless channels could use duplicate signals over satellite communications, providing multiple pathways to ensure control that have largely non-overlapping vulnerabilities. Interestingly, another approach might be to partially revert back to old technology such as SCADA implemented using proprietary protocols. This would require attackers to have detailed information about each system, rather than exploiting vulnerabilities in general-purpose support systems. This also makes it possible to use knowledge of particular details of components in the system as sources for additional sensing to corroborate conventional awareness measures.

For this vision to truly achieve cyber resilience for cyber physical systems, more information sharing will be required between organizations and critical infrastructure owners and operators. This will require engagement with the National Council of ISACs [Information Sharing and Analysis Centers] in a way that results in common language and multi critical infrastructure communications.

Specific proposed research steps include:

1) Solve the primary economic issue. Accurately valuate energy and find a way to share cyber resilience costs via increased rates, grants, or other means so that utility owners need not cover the entire cost of achieving cyber resilience.
2) Convince public utility commissions and customers that cyber resilience in cyber physical systems is useful and worthwhile to invest in.
3) Incentivize or issue new requirements to manufacturers
    a. Enforce NIST 800-53.
    b. Add new capabilities to existing systems.
    c. Increase capacity of systems to handle the additional load of resilience technologies.
4) Implement ubiquitous, dynamic, remote patching.
    a. Ensure that this does not introduce new cyber vulnerabilities.
5) Create energy storage/sinks as a service redundancy and if possible a symmetric distribution/collection network.
    a. Harness multiple types of energy sources, although this introduces other challenges.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

b. Implement collection and management systems for solar, wind, and other alternative energy sources.

Findings and Recommendations

One major finding is that resilience needs for cyber physical systems can be highly specific for the different critical infrastructure domains because of the highly specialized tasks performed by the physical components in these systems, but some aspects of resilience still transcend these domains. One such issue is the need for systems to continue functioning in the face of failures within the system and even despite failures in highly interconnected supporting critical infrastructure domains in which resilience response may lie entirely outside the scope of influence of the cyber physical system of interest. Communications infrastructure, GPS, and other ubiquitous services support multiple other critical infrastructure domains and therefore drive the need for alternate out-of-band methods for realizing redundancy in the event of failure in one of these domains. This will require concerted new research efforts into cross-critical infrastructure domain modeling and simulation as well as technology solutions for redundant pathways to deliver some of the services currently provided by a single technology.

The group also identified redundancy and diversity as the two primary means by which resilience could be achieved in cyber physical systems. It remains to be seen whether these two methods can be implemented in economically viable solutions, and whether these will, in fact, confer the desired level of resilience. To evaluate these assertions, new economic models will be required along with rigorous measures, testing protocols, and threat models that truly reflect the cyber environment in which cyber physical systems must operate. The group also identified domain-specific issues for the representative domains of energy and transportation. Many of these have applicability to other domains, but is the group left it to future meetings to explore cyber physical systems resiliency needs in other domains specifically.

In conclusion, resilience is widely considered a desirable property for critical infrastructure. But as these domains increasingly rely on both cyber and physical components, connected in geographically distributed and complex ways, the cyber physical aspects of these systems will require their own resilience-related research and development programs. Ultimately the goal is to ensure that the processes essential to supporting the U.S. way of life persist in the face of growing, constant cyber conflict.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

6. Cyber Resilience Tabletop

Track Chair: Harriet Goldman, The MITRE Corporation
Track Scribe: Ellen Laderman, The MITRE Corporation

6.1 Goal

Cyber resiliency tabletop exercises provide opportunities for participants to enhance their understanding of how cyber resiliency can mitigate the impact of cyber adversaries and the benefits and trade-offs of implementing cyber resiliency. The goals of this track were to gain insights about how cyber resiliency tabletop exercises can be more effective, and to enhance participants' understanding of:

- How a determined cyber adversary creates different challenges from those addressed by current forms of operational resilience or continuity of operations (COOP) plans and disaster recovery plans
- How determined cyber adversaries approach defeating protections and plan their exploits in order to achieve their intended outcomes
- How to weigh the benefits and impacts offered by selected cyber resiliency techniques (e.g., synergies, dependencies, and potential conflicts among the techniques)
- How adopting a cyber resiliency mindset aligns with a risk management approach and differs markedly from a compliance mindset
- How to strike a balance between proactive architectural techniques and responsive actions

6.2 Discussion / Observations

The tabletop exercise scenario was based on a fictional logistics company, modelled roughly after FedEx, UPS, the USPS, and a variety of other national and regional carriers. The logistics company in the exercise owns a large fleet of trucks that use transshipment facilities in numerous locations to move packages from their sources to their destinations. Company operations rely on a dynamically updated distributed database that tracks packages and trucks. Handheld scanners are used throughout the shipment process: when packages are picked up, transferred from a truck to a facility or from a facility to a truck, and delivered to their final destinations. The company also uses third-party suppliers for services outside its core business expertise (e.g., HVAC and maintenance).

Due to the limited time available, this tabletop exercise did not include actual systems or a mock environment. The track participants were divided into two Blue Teams (i.e., defenders) with White Team members (i.e., control to guide the exercise if needed) for each Blue Team. The Red Team (i.e., the adversary) was represented by six inject points where the Blue Team was informed of adverse impacts, as shown in Figure 5. The scope of this exercise covered the cyber physical systems and the enterprise information and communication technology, with a focus on the company's roles critical to securing, operating, and defending the business. Key considerations during the exercise included:

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

- What resilience options would the Blue Team consider?
- What technical solutions are most applicable?
- What changes in policies, procedures, or governance are needed?
- What constraints does the Blue Team need to take into account – e.g., cost factors, operational limitation, or policy barriers?

Defender roles (Blue Team) included system administrator, front-line defenders in the company's computer security operations center (CSOC), CSOC manager, IT managers or business process owners, and security architects.

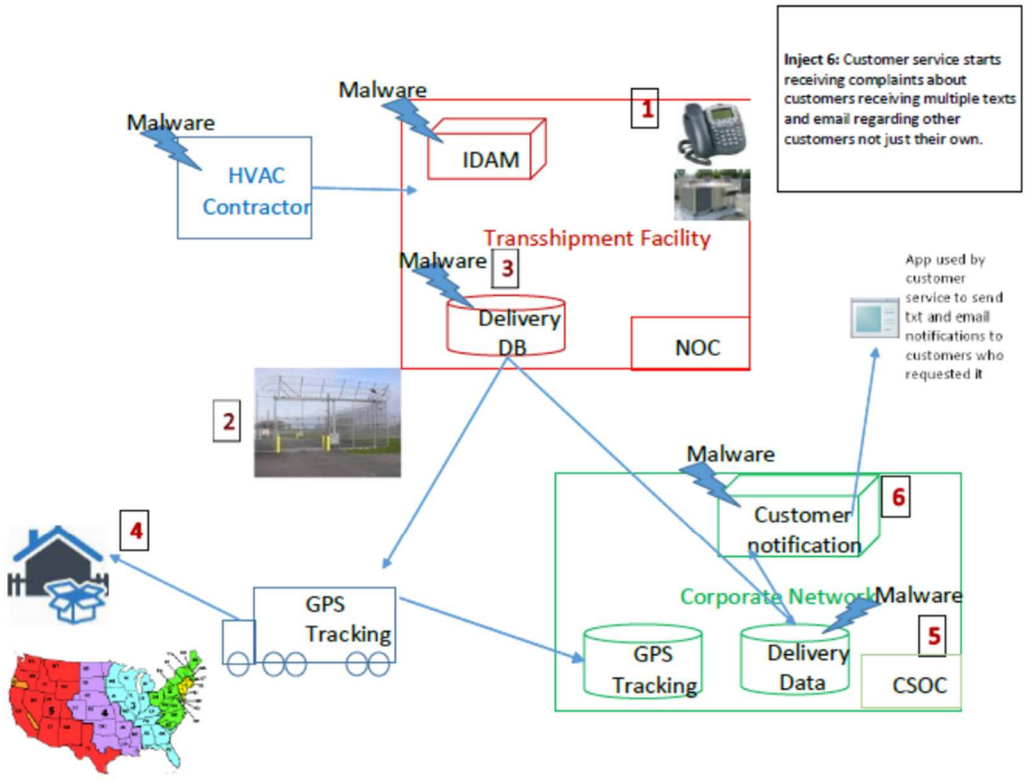

Figure 5. Exercise Scenario and the Six Injects

The track participants received a brief overview of the cyber resiliency framework [6] so that they had a common language and background on techniques that could be applied during the exercise.

After each inject (i.e., adverse event), the Blue Teams had the opportunity to select cyber resiliency techniques to apply to the environment and situations.

6.3 Challenges

Similar to what would happen in a real-world attack, participants were initially unclear about what was occurring and the meaning of events. Two main challenges faced by the participants were:

- Understanding the importance of the company's critical business objectives in the context of what was occurring
- Understanding when to escalate incident reporting and response

The first challenge pointed out the difference between mission objectives and local concerns such as keeping systems up and functioning. At the start of the exercise the participants focused on protecting individual systems. This changed as the exercise progressed and more systems were attacked.

The second challenge, understanding the timing for incident reporting escalation, pointed to the need for specific cyber resiliency implementation, particularly with respect to situational awareness and how the adverse circumstances affected the business mission objectives. As the exercise progressed this need became more clear and this generated discussion about choices made earlier in the exercise and how they would be changed if time could be rewound.

The participants found the *Cyber Resiliency Engineering Aid* [6] useful in creating a context for the engagement. They acknowledged the many ways that this problem can be framed and agreed that the *Engineering Aid* covered most aspects effectively. The hands-on application through the exercise proved very useful in understanding the techniques and objectives discussed in the *Engineering Aid.*

Summary

The challenges posed by a determined cyber adversary (e.g., a stealthy, persistent, and sophisticated adversary, who may have already compromised system components and established a foothold within an organization's systems) differ from the challenges addressed in current COOP and Disaster Recovery Plans (which postulate an easily recognized adverse event or set of adverse conditions). The tabletop exercise illustrated the difference between protecting the mission/business and protecting information technology. Protecting the mission/business is driven by resiliency goals of withstanding (i.e., keeping the business going in adverse circumstances), recovering from adversity, and evolving so that in the future cyber adversaries have less impact. All of this is predicated on preparing properly (e.g., having a clear understanding of how information and communications technology supports the mission).

Stakeholder engagement is crucial to the effectiveness of resilience-oriented analyses. In order to communicate well a common framework or model is critical. Such common models or frameworks help stakeholders make sense of the problem and solution domains together.

6.4 Recommendations/Way Forward

After completing the exercise, the participants had six recommendations. Three of them related to conducting cyber resiliency assessments and how to make them most effective, and three related to making tabletop exercises useful in cyber resiliency assessments. The three recommendations about cyber resiliency assessments are:

- Engage all stakeholder types (e.g., legal, information technology, business operations, policy, public relations, information security) across the organizations involved in mission success.
- Remember that an assessment is not just an IT or engineering exercise; it is about the mission.
- Include local subject matter experts, management, and the work force (boots on the ground) – it is important to get both breadth and depth.

The three recommendations about tabletop exercises and follow-on activities to make them useful in cyber resiliency assessments are:

- Keep it fun! – Because this activity was enjoyable the participants stayed engaged and learned more than they would have otherwise. The two Blue Teams did not compete and so participants were open to listening to each other.
- The mission is why the system exists in the first place. Starting with well-defined mission threads ensures that the use cases cover the essential functionality and environment.
- Expand use cases to develop exercises and simulations. This helps ensure that the use cases are realistic and systems are really prepared for the mission.

This report was compiled from sessions at the MITRE Fifth Annual Secure and Resilient Cyber Architectures Invitational, 2015
Approved for Public Release. Distribution Unlimited. Case Number: 16-1199

7. References

[1] The MITRE Corporation (ed.), "2nd Secure and Resilient Cyber Architectures Workshop: Final Report," 2012. [Online]. Available: https://registerdev1.mitre.org/sr/2012_resiliency_workshop_report.pdf.

[2] The MITRE Corporation (ed.), "Third Annual Secure and Resilient Cyber Architectures Workshop," December 2013. [Online]. Available: http://www.mitre.org/sites/default/files/publications/13-4210.pdf.

[3] The MITRE Corporation (ed.), "Fourth Annual Secure and Resilient Cyber Architectures Invitational," 2015. [Online]. Available: http://www.mitre.org/sites/default/files/pdf/2014-Secure-Resilient-Cyber-Architectures-Report-15-0704.pdf.

[4] J. Tighe (VADM), "Statement of Vice Admiral Jan Tighe Before the Subcommitte on Emerging Threats and Capabilities of the House Armed Services Committee," 4 March 2015. [Online]. Available: http://docs.house.gov/meetings/AS/AS26/20150304/103093/HHRG-114-AS26-Wstate-TigheJ-20150304.pdf.

[5] CERT Program, "CERT® Resilience Management Model, Version 1.0: Improving Operational Resilience Processes," May 2010. [Online]. Available: http://www.cert.org/archive/pdf/10tr012.pdf. [Accessed 26 October 2011].

[6] D. Bodeau, R. Graubart, W. Heinbockel and E. Laderman, "Cyber Resiliency Engineering Aid-The Updated Cyber Resiliency Engineering Framework and Guidance on Applying Cyber Resiliency Techniques, MTR140499R1, PR 15-1334," May 2015. [Online]. Available: http://www.mitre.org/sites/default/files/publications/pr-15-1334-cyber-resiliency-engineering-aid-framework-update.pdf or http://www.defenseinnovationmarketplace.mil/resources/20150527_Cyber_Resiliency_Engineering_Aid-Cyber_Resiliency_Techniques.pdf.

[7] NIST, "NIST SP 800-39, Managing Information Security Risk: Organization, Mission, and Information System View," March 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-39/SP800-39-final.pdf.

[8] D. J. Snowden and M. E. Boone, "A Leader's Framework for Decision Making," *Harvard Business Review*, November 2007.

[9] G. Gandhi, "Complexity in cyber security," BCS: The chartered institute for IT, September 2015. [Online]. Available: http://www.bcs.org/content/conWebDoc/55148.

[10] S. Ali, "Crisis measures to improve cyber resilience a 7 step approach," LinkedIn, 20 March 2015. [Online]. Available: https://www.linkedin.com/pulse/crisis-measures-improve-cyber-resilience-7-step-approach-sheraz-ali.

[11] International Risk Governance Council, "An Introduction to the IRGC Risk Governance Framework," February 2008. [Online]. Available: http://irgc.org/wpcontent/uploads/2012/04/An_introduction_to_the_IRGC_Risk_Governance_Framework.pdf.

[12] Evidence Based Research, Inc., "Network Centric Operations Conceptual Framework Version 2.0 (DRAFT): A Collaborative Effort of John Gartska, Office of Force Transformation & David

lberts, Office of Assistant Secretary of Defense (Networks and Information Integration)," June 2004. [Online]. Available: http://www.ibrarian.net/navon/paper/Network_Centric_Operations_Conceptual_Framework_V.pdf?paperid=4455264.

[13] NIST, "Framework for Improving Critical Infrastructure Security, Version 1.0," 12 February 2014. [Online]. Available: http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214.pdf.

[14] D. Bodeau, Graubart and Richard, "Cyber Resiliency and NIST Special Publication 800-53 Rev.4 Controls (MTR 130531, PR 13-4037)," September 2013. [Online]. Available: http://www.mitre.org/sites/default/files/publications/13-4047.pdf.

[15] NIST, "Federal Information Processing Standards Publication (FIPS PUB) 199: Standards for Security Categorization of Federal Information and Information Systems," February 2004. [Online]. Available: http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf.

[16] NIST, "Security and Privacy Controls for Federal Information Systems and Organizations (NIST SP 800-53 R4)," April 2013. [Online]. Available: http://dx.doi.org/10.6028/NIST.SP.800-53r4.

[17] NIST, "Guide for Conducting Risk Assessments, NIST SP 800-30 Rev.1," September 2012. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-30-rev1/sp800_30_r1.pdf.

[18] CNSS, "Security Categorization and Control Selection for National Security Systems (CNSSI No. 1253), Version 2," 15 March 2012. [Online]. Available: http://www.disa.mil/Services/DoD-Cloud-Broker/~/media/Files/DISA/Services/Cloud-Broker/cnssi-security-categorization.pdf.

[19] DoD Defense Science Board, "Task Force Report: Resilient Military Systems and the Advanced Cyber Threat," January 2013. [Online]. Available: http://www.acq.osd.mil/dsb/reports/ResilientMilitarySystems.CyberThreat.pdf.

[20] D. Bodeau and R. Graubart, "Cyber Resiliency Engineering Framework (MTR110237, PR 11-4436)," September 2011. [Online]. Available: http://www.mitre.org/sites/default/files/pdf/11_4436.pdf.