Third Annual Secure and Resilient Cyber Architectures Workshop

Overview

The Third Annual Secure and Resilient Cyber Architectures Workshop, cooperatively supported by The MITRE Corporation, the National Security Agency (NSA) R2, and the NSA Information Assurance Directorate, took place on June 19–20, 2013, at the MITRE facilities in McLean, VA. The NSA co-sponsors played an important role in determining the objectives, expected outcomes, and content. In addition, NSA personnel led two of the workshop tracks. Dr. Ron Ross of the National Institute of Standards and Technology played a major role in shaping the workshop and led one of the tracks.

The workshop drew substantially more attendees than the previous year's event: Day 1, featuring keynotes and briefings by government leaders, attracted approximately 100 attendees; Day 2, which featured three simultaneous tracks focused on key issues, attracted 60–65 attendees. This report presents the results of the discussions and follow-on interactions during the topical tracks—Using a Risk-based Approach to Select and Apply Cyber Resiliency Techniques, Impacting Planned Architectures for the Future, and Advancing Cyber Resiliency Through Active and Adaptive Response. The report captures the points that participants considered the most salient portions of the discussion, items of consensus, questions raised during the discussions, and comments on next steps. Participants in each track reviewed these summaries through email exchanges. All other materials from the workshop, as well as the agenda and briefings, can be found at https://register.mitre.org/sr/materials.html and https://register.mitre.org/sr/agenda.html.

We welcome comments from readers through the contact email address: secureandresilient@mitre.org.

The Cyber Resiliency Workshop Committee

November 2013

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013

Executive Summary

The Third Annual Secure and Resilient Cyber Architectures Workshop was held June 19–20, 2013. The purpose of the workshop was to bring together a diverse group of experts to advance cyber resiliency concepts and develop guidance for architecture, design, and research and development investment. Day 1 featured keynote addresses by government and commercial leaders. The briefings that day included a threat perspective view of cyber resiliency, various cyber resiliency use cases, cyber resiliency case studies, and a view of cyber resiliency from various leading commercial organizations.

Day 2 consisted of facilitated working groups in three tracks:

- Track 1: Using a Risk-based Approach to Select and Apply Cyber Resiliency Techniques
- Track 2: Impacting Planned Architectures for the Future
- Track 3: Advancing Cyber Resiliency Through Active and Adaptive Response.

Track 1

Track 1 focused on the use of existing risk-based guidelines and frameworks, with the goal of producing consensus guidance on best practices for a variety of resiliency techniques. Three underlying assumptions drove the discussions:

- There is no single best resiliency technique to apply in all systems/environments.
- There is no minimum set of resiliency techniques to be applied.
- The selection of an optimum set of techniques depends on various risk factors.

Based upon the group discussions, the participants proposed an initial set of risk factors for consideration for determining which resiliency techniques to employ:

- Relevance of techniques
- Organizational goals and objectives
- Effectiveness of the techniques in addressing threats of concern
- Maturity of techniques
- Operational application of the techniques in current practice
- Political, operational, economic, and technical (POET) considerations in employing the techniques
- Feasibility of applying the techniques
- Capability, intent, and targeting of adversary against whom the techniques are to be applied
- Stakeholder buy-in.

Track 1 participants identified six follow-on actions. The participants recognized that while some could be done in the near term, others could take many years.

- Map the National Institute of Technology Special Publication (NIST SP) 800-53 security controls to resiliency techniques.

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013
Approved for Public Release. Case Number 13-4210. Distribution Unlimited.

- Build a resiliency overlay that would identify those security controls needed to ensure a mission's cyber resiliency even in the presence of the advanced persistent threat (APT).
- Explore development of additional resiliency techniques to ensure completeness of the set of techniques.
- Develop awareness briefings for mission owners to ensure that they are involved in and supportive of the integration of cyber resiliency into their missions.
- Create more prototypes and experiments to validate the effectiveness of techniques.
- Ascertain and collect the cost of the various resiliency techniques.

Track 2

Track 2 focused on developing a prioritized list of cyber resiliency techniques that would be most beneficial within the future Department of Defense (DoD)/Intelligence Community (IC) information technology environments—Joint Information Environment (JIE) and IC Information Technology Enterprise (ITE). The participants developed five high-level principles for the DoD/IC to consider in adopting cyber resiliency techniques and a notional prioritization of the techniques that would be most beneficial if applied now within the target architectures.

- All of the techniques would be useful and beneficial; prioritization based upon various considerations will be required.
- Some of the techniques simply represent good engineering practice and are already reflected in the JIE and ITE architectures.
- Some techniques, if not explicitly addressed early in the system life cycle, will be executed poorly.
- Careful scoping and coordination between the mission and enterprise communities will be necessary to incorporate cyber resiliency effectively.
- Cost-benefit trade-offs must be examined to identify best fit techniques and calculate cyber resiliency return on investment.

Track 3

Track 3 focused on methods for applying static and dynamic response mechanisms that quantitatively improve cyber resiliency and measurably affect adversary behavior. The attendees were divided into teams to develop their own realistic cyber resiliency narratives or vignettes, focusing on one of five predefined areas: the "Internet of things," the mobile workforce, cyber-enabled transportation, safety-critical systems, and symbiotic systems. Each vignette consisted of a description of the particular system or scenario and insight into the challenge of maintaining system resiliency in the presence of persistent threats.

Three major themes were consistently brought up during discussion and were common across all vignettes:

- As the Internet becomes more prevalent and mature, the physical and cyber worlds become more intertwined. Cyber resiliency is critical to protect and enable operation of "smart" and cyber-enabled products.

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013

- Time is a critical factor: it is necessary to monitor changes and relative rates of change, or measures response effectiveness in terms of the time to react or resolve the issue.
- Enhancing resiliency will require collaboration across many organizations and will place a high value on establishing trust, authorization, and reliable cross-domain data sharing.

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013

Table of Contents

Overview i

Track 1 iii

Track 2 iv

Track 3 iv

Track 1: Using a Risk-based Approach to Select and Apply Cyber Resiliency Techniques 1

Background 1

Objective 1

Challenges 1

Agreed-Upon Resiliency Risk Factors 3

Cyber Resiliency Guidance Needs 6

Additional Topics for Consideration 8

Follow-on Actions 8

Track 2: Impacting Planned Architectures for the Future 9

Applying Cyber Resiliency to Planned Architectures 9

Cyber Resiliency Techniques 10

Adaptive Response 10

Analytic Monitoring 10

Coordinated Defense 11

Deception 11

Diversity 11

Dynamic Positioning 11

Dynamic Representation 12

Non-persistence 12

Privilege Restriction 12

Realignment 12

Redundancy 13

Segmentation 13

Substantiated Integrity 13
Unpredictability 14
Session Outcomes 14
Track 3: Advancing Cyber Resiliency Through Active and Adaptive Response 17
Foundations for Cyber Resilience 17
Resilient Response Framework 18
Working Sessions—"Future of Resilience" 20
Cyber Adversary 21
"Internet of Things" 23
Mobile Workforce 24
Cyber-Enabled Transportation 25
Safety-Critical Systems 26
Symbiotic Systems 28
Notable Findings 30
Value Proposition 32

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013
Approved for Public Release. Case Number 13-4210. Distribution Unlimited.

List of Figures

Figure 1. MITRE's Cyber Resiliency Engineering Framework 2
Figure 2. Maturity vs. Operational Application of Cyber Resiliency Techniques 3
Figure 3. Differing Directions of Technologists and Mission Owners 6
Figure 4. Prioritized Techniques 15
Figure 5. Response Matrix 21

List of Tables

Table 1. Sample POET Considerations and Their Impact on Cyber Resiliency Techniques 4
Table 2. Presented Working Definitions 18
Table 3. Cyber Adversary Response Matrix 22
Table 4. "Internet of Things" Response Matrix 24
Table 5. Mobile Workforce Response Matrix 25
Table 6. Cyber-Enabled Transportation Response Matrix 26
Table 7. Safety-Critical Systems Response Matrix 27
Table 8. Symbiotic Systems Response Matrix 29

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013
Approved for Public Release. Case Number 13-4210. Distribution Unlimited.

Track 1: Using a Risk-based Approach to Select and Apply Cyber Resiliency Techniques

Track Chair: Ron Ross, National Institute of Standards and Technology (NIST); ron.ross@nist.gov
Co-chair: Richard Graubart, The MITRE Corporation; rdg@mitre.org

Background

This track focused on the use of existing risk-based guidelines and frameworks, with the goal of producing consensus guidance on best practices for a given resiliency technique. Three underlying assumptions drove the discussions:

- There is no single best resiliency technique to apply to an information system.
- There is no minimum set of resiliency techniques to be applied.
- The choice of the optimum set of techniques depends on various risk factors.

In short, with regard to cyber resiliency, one size does not fit all. Given those assumptions, the discussions during this track aimed at helping organizations to determine the optimum resiliency techniques for their purposes.

Objective

Because no single best set of cyber resiliency techniques (and associated National Institute of Standards and Technology Special Publication [NIST SP] 800-53 controls) exists, the question arises: What process can organizations use to identify the resiliency techniques most appropriate to meeting their particular needs? Toward that end, the working group proposed the following objectives:

- Identify the risk factors to consider in determining appropriate resiliency techniques.
- Categorize risk factors by:
  - Relative importance
  - Maturity (i.e., how well established/researched).
- As feasible, identify less established risk factors.
- Identify gaps in guidance requiring further study.

Moreover, because the workshop sought to produce actionable recommendations, the participants hoped to achieve sufficient consensus to enable progress toward establishing cyber resiliency guidelines (e.g., a cyber resiliency overlay).

Challenges

The lack of a well-defined operational risk management framework poses a challenge for defenders. NIST SP 800-39, Managing Information Security Risk, produced jointly by NIST, the Department of Defense (DoD), and the Intelligence Community (IC), only contains guidance at a

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013

very high level. Transforming the broad tenets of that document into actionable processes for determining appropriate cyber resiliency techniques remains an ongoing process.

To provide a common vehicle for discussion, the meeting began with quick overview of the MITRE Cyber Resiliency Engineering Framework. Figure 1 depicts the framework, with its four goals, eight objectives, and 14 techniques. The techniques appropriate for one organization may differ greatly from those appropriate for another. In addition, it would not be financially feasible for all organization to apply the same techniques.

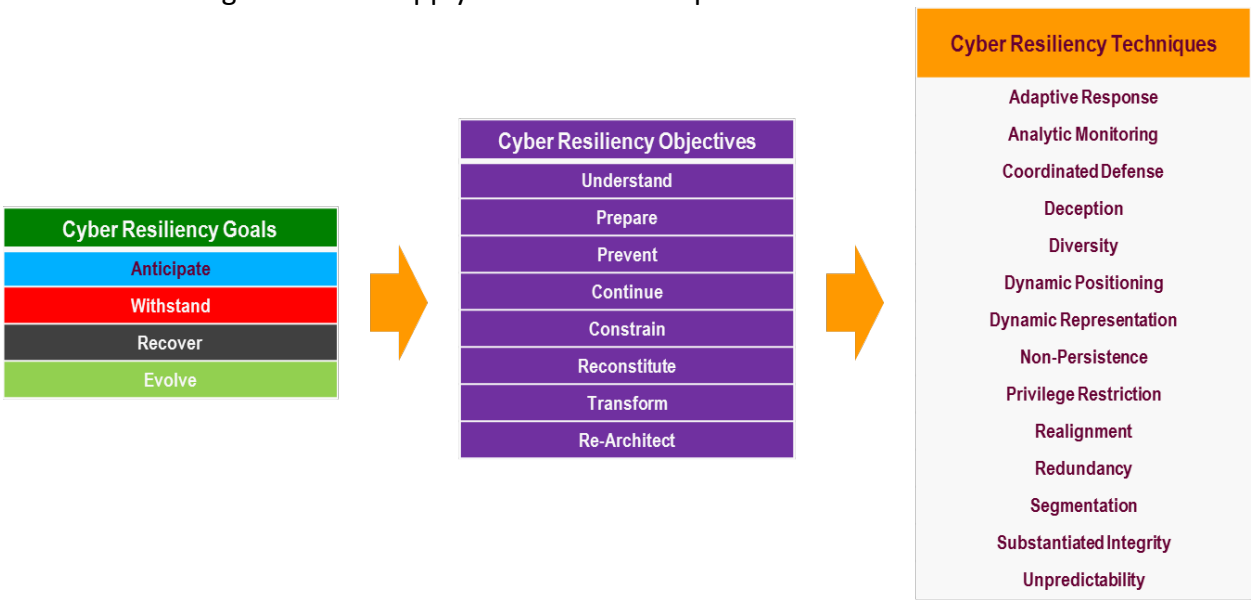

Figure 1. MITRE's Cyber Resiliency Engineering Framework

Discussants pointed out the wide variation in the relative maturity and adoption of the various techniques (see Figure 2). They also noted that the issue has a third dimension (not reflected in Figure 2): the relative effectiveness of the techniques. Just because a technique is mature or in common use in the community does not mean it is highly effective, especially against the advanced persistent threat (APT).

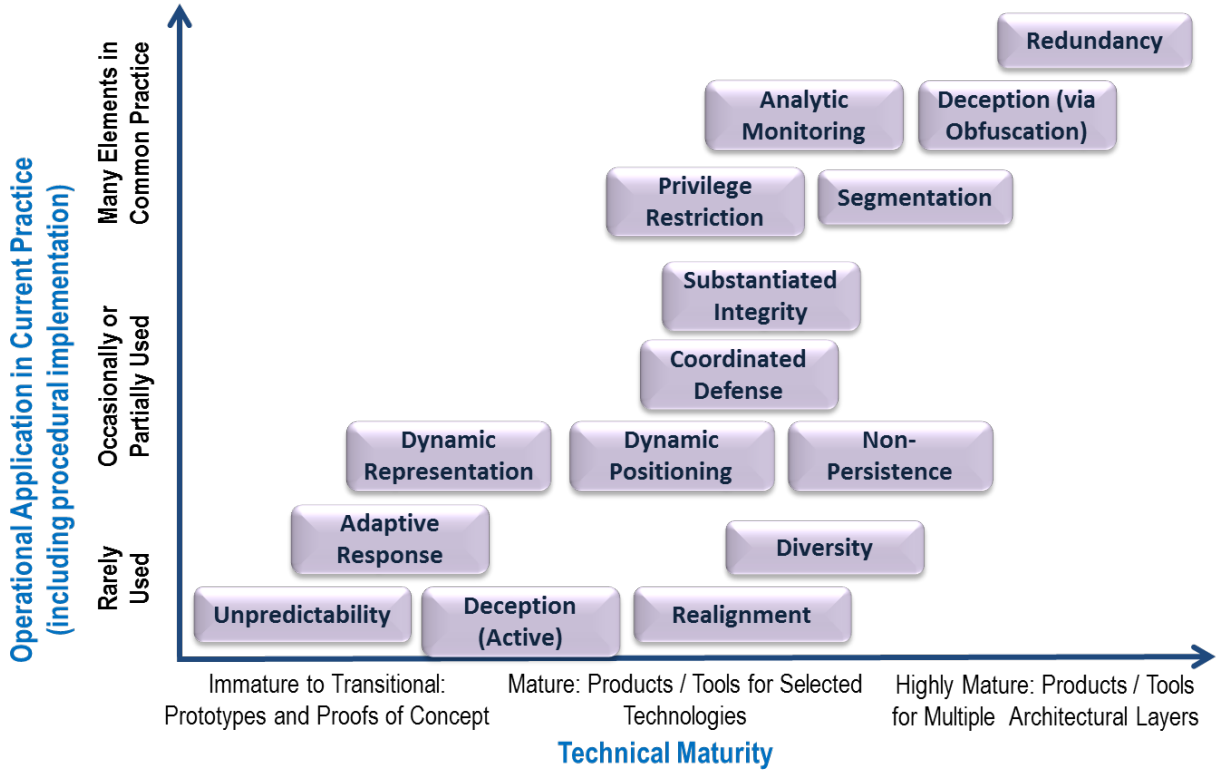

Figure 2. Maturity vs. Operational Application of Cyber Resiliency Techniques

Opinions regarding the relative importance of a given technique, and even the factors used to assess a given technique, vary according to organizational perspective. For example, personnel involved in acquisition would probably place high importance on the relative maturity of a technique, but personnel engaged in cyber security operations would likely base their choice primarily on the relative effectiveness of the technique.

The discussion also pointed out that some techniques can be automated, while others require user interaction. For those requiring user interaction, organizations must also consider the relative expertise of the user/operator. Some techniques may demand greater expertise than most operators possess, thus confusing the operators and potentially impeding the effectiveness of the techniques.

Agreed-Upon Resiliency Risk Factors

On the basis of the discussions, the participants proposed an initial set of risk factors for consideration:

- **Relevance of techniques**: Not all cyber resiliency techniques are equally relevant to an organization. The organization's mission and/or operating environment may impose

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013
Approved for Public Release. Case Number 13-4210. Distribution Unlimited.

constraints on the viability and utility of certain techniques. For example, environments with very limited processing capability (e.g., embedded systems) would likely not find techniques such as deception nets very useful.

- **Organizational goals and objectives**: Different stakeholders have different goals and objectives. Mission commanders focus primarily on ensuring full execution of the mission; thus, they may be most interested in the goals of Anticipate and Withstand (see Figure 1) and the objectives and techniques that support those goals. By contrast, cyber defenders probably emphasize the ability to respond quickly to cyber attacks in their operational environment. Given such concerns, these stakeholders might accord special importance to the goal of Evolve (and the supporting objectives and techniques) in order to obtain the optimum tools to respond to the changing threat posed by the APT.
- **Effectiveness in addressing threats of concern**: Ideally, the community would have some means of quantifying the effectiveness of a resiliency technique, both relative to other techniques and against specific classes of threats. Further, because resiliency techniques do not operate in a vacuum, the community should assess the effectiveness of the techniques relative to the cost of applying them.
- **Maturity of techniques**: As noted in Figure 2, the maturity of techniques varies. For some stakeholders, the relative maturity of a given technique is important in determining whether or not to invest in and select it.
- **Operational application in current practice**: Again as noted in Figure 2, some techniques are used more commonly than others. Many organizations may select popular techniques because they (sometime mistakenly) equate common practice with best practice. But simply because a technique has a proven track record in some environments against a particular set of threats does not mean it will prove effective against a different set of threats. For example, physically distributed redundant copies of systems offer effective protection in the event of fire or flood, but provide no safeguards against malware targeted to the systems that the organization employs.
- **Political, operational, economic, and technical (POET) considerations**: Various POET considerations can greatly influence whether an organization might select a given cyber resiliency technique. Table 1 shows a representative set of such considerations.

Table 1. Sample POET Considerations and Their Impact on Cyber Resiliency Techniques

Technique	Representative Reasons for Restricting Consideration
Adaptive Response	Liability concerns (e.g., responses that violate service-level agreements, cause collateral damage)
Analytic Monitoring	Policy concerns related to collecting, aggregating, and retaining data (e.g., sensitivity/classification, privacy)
Coordinated	Governance and concept of operations (CONOPS) issues (e.g., overlapping or

Technique	Representative Reasons for Restricting Consideration
Defense	incompletely defined roles and responsibilities; no clear responsibility for defining cyber courses of action [COAs])
Deception	• Legal, regulatory, contractual, or policy restrictions • Concern for reputation
Diversity	• Policy or programmatic restrictions (e.g., organizational commitment to a specific product or product suite) • Life-cycle cost of developing or acquiring, operating, and maintaining multiple distinct instances
Dynamic Positioning	Technical limitations due to policy or programmatic restrictions (e.g., organizational commitment to a specific product or product suite that does not accommodate repositioning)
Dynamic Representation	Governance issues/information sharing constraints in the context of systems-of-systems
Non-persistence	Technical limitations that prevent refresh functions from meeting quality of service (QoS) requirements
Privilege Restriction	Governance and CONOPS issues (e.g., inconsistencies or gaps in definitions of roles, responsibilities, and related privileges; operational impetus to share roles)
Realignment	Organizational and cultural impacts (e.g., eliminating functions that personnel are used to employing, impact on morale of relocating staff)
Redundancy	Costs of maintaining multiple, up-to-date, and secure instantiations of data and services
Segmentation	Cost and schedule impacts of re-architecting; cost of additional routers, firewalls
Substantiated Integrity	Cost and schedule impacts (e.g., of incorporating and managing cryptographic checksums on data)
Unpredictability	Operational and cultural issues (e.g., adverse impact on planned activities, adverse impact on staff expectations of how to operate)

- **Feasibility of applying the techniques**: Organizations/stakeholders must be able to employ the technique and achieve a concrete benefit from it.
- **Capability, intent, and targeting of adversary**: Capability, intent, and targeting refer to the characteristics that defenders examine to assess an adversarial threat. These characteristics may have a direct bearing on determining which resiliency techniques are most appropriate to counter a particular adversary. For example, defenders might best counter an adversary capable of developing and deploying new zero day attacks tailored to the systems that the defender is known to use by employing techniques such as deception (e.g., deception nets). This technique would allow defenders to discern the nature of the attack techniques and diversify (e.g., use different systems, services, or applications at key locations) and devise methods to impede the adversary's ability to target the organization's systems.

- **Stakeholder buy-in**: Many participants noted that it was not sufficient to merely identify and develop cyber resiliency techniques: the mission owners must become involved and supportive as well. Mission owners and security technologists often talk past each other, like two trains on parallel tracks (see Figure 3). Participants agreed on the importance of explaining the security threats and relative advantages and disadvantages of the various cyber resiliency mitigations/controls to the mission owners in terms that they understand in order to obtain their support. The mission owner may not care about the technical details and nuances of the various resiliency techniques; in other words, the "how" of a technique is not necessarily relevant. However, the mission owner needs to understand "what must be done" and "why it must be done." Cyber defenders must explain the techniques chosen to the mission stakeholders in the context of:
  - Operational effectiveness
  - Suitability
  - Survivability.

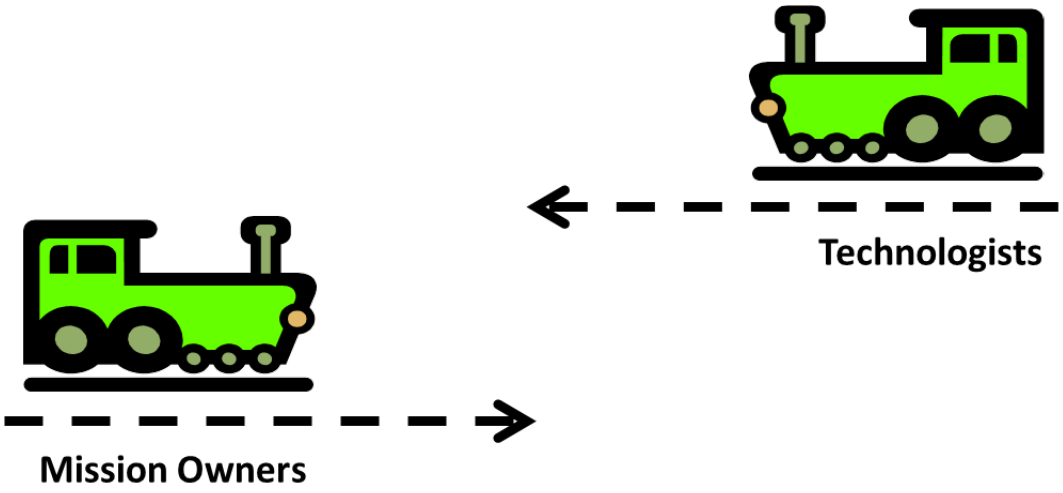

Figure 3. Differing Directions of Technologists and Mission Owners

Cyber Resiliency Guidance Needs

Participants recognized the need for certain guidance documents to meet the objectives of this track.

- **Identifying NIST SP 800-53 security controls that address resiliency**: NIST SP 800-53 presents more than 860 controls and enhancements. A very large number address various aspects of cyber resiliency (several dozen controls were added specifically for that purpose). However, the sheer number of controls makes it difficult for those developing system requirements as well as overlays to identify which controls support

resiliency and, of those, which resiliency techniques they support. The explicit identification of the controls that support resiliency would be beneficial.

- **Resiliency overlay**: An overlay is a specification of security controls, control enhancements, supplemental guidance, and other supporting information employed during the tailoring process that complements (and further refines) security control baselines. The overlay specification may be more or less stringent than the original security control baseline specification and can be applied to multiple information systems. The NIST baselines explicitly state that they do not assume the presence of the APT. A resiliency overlay would complement the baselines by identifying those security controls intended to ensure a mission's cyber resiliency even in the presence of the APT. Such an overlay would likely encompass several component overlays, reflecting the different nuances (e.g., objectives) that drive the selection of mitigations.
- **Handbook(s) that express resiliency "requirements" in the language of the audience**: As mentioned earlier, defenders must convey the key aspects of cyber resiliency to stakeholders (e.g., mission owners) in their own language. Different handbooks would likely be required to address the needs of different types of stakeholders.
- **Integration/mapping into/to larger national/international standards (crosswalks of standards)**: While incorporating cyber resiliency into government guidance (e.g., NIST Special Publications) is valuable, it is not necessarily sufficient. Many organizations, especially commercial organizations whose customer base extends beyond the government, must ensure that cyber resiliency concepts are integrated into national and international standards.
- **Mapping of threats against resiliency techniques/controls**: It is not sufficient merely to identify cyber resiliency controls. Some controls are more effective against certain cyber threats than others; some cyber threat techniques occur at different stages of the cyber attack life cycle (aka the Kill Chain); and some cyber controls are more effective at some stages of the life cycle than others. Therefore, defenders need a mapping of cyber resiliency controls to different cyber threats. However, this mapping must be preceded by the development of a national cyber threat database at the unclassified level. To be useful, such a database must be frequently updated to ensure that the spectrum of threats remains current. Similarly, any mapping of cyber resiliency controls against cyber threats depends on constant updates to remain current.
- **Special publication on cyber resiliency**: Defenders need a special publication that presents a clear, authoritative, understandable discussion of cyber resiliency. In addition to explaining the major constructs of cyber resiliency, such a document would likely contain or at least point to cyber resiliency overlays and any relevant resiliency controls described in SP 800-53.

Additional Topics for Consideration

Given the limited time available during the workshop, the participants could examine only the topics discussed above. However, the workshop did devote some time to identifying important related topics not addressed during the track.

- **Training and awareness activities to affect the culture**: Resiliency is still an unfamiliar concept for most users. Simulations and exercises could help improve awareness of the benefits and limitations of cyber resiliency. For example, training programs could place operators in a simulated cyber attack scenario where they could employ various cyber resiliency techniques to help them better understand which techniques work best against which attacks.
- **Completeness of techniques**: The 14 techniques of the MITRE Cyber Resiliency Engineering Framework cover a very broad set of possible mitigations, but even so the framework may not be exhaustive.
- **Effect of implementing resiliency techniques on mission operations**: Like any new technique or approach, cyber resiliency techniques may have unexpected impacts on operations. These could range from a need for more training of operators to unanticipated costs for maintaining some techniques. As organizations begin to adopt some of the techniques, the effects will become clearer.

Follow-on Actions

Track participants identified six follow-on actions. Some of these can be performed in the near term (e.g., mapping the NIST SP 800-53 controls to resiliency techniques), but others (e.g., collecting cost data) may take years.

1. Map NIST SP 800-53 controls to resiliency techniques.
2. Build resiliency overlay(s).
3. Explore development of additional resiliency techniques.
4. Develop awareness briefings for mission owners.
5. Create more prototypes and experiments to validate the effectiveness of techniques.
6. Ascertain and collect the cost of the various resiliency techniques.

Track 2: Impacting Planned Architectures for the Future

Track Chair: Bryan Larish, National Security Agency (NSA)
Co-chair: Roger Westman, The MITRE Corporation; rwestman@mitre.org

The second track focused on developing strategies for adopting cyber resiliency techniques in future DoD/IC information technology environments. The track began with an introduction by NSA and a high-level view of MITRE's Cyber Resiliency Engineering Framework presented by MITRE. MITRE highlighted key cyber resiliency goals and explained each of the 14 cyber resiliency techniques that would serve as the focus of the latter part of the session. The subsequent discussion centered on the availability/maturity of technology that implements the various techniques and how the techniques align with the cyber attack life cycle. Following that presentation, subject matter experts from MITRE briefed overviews of DoD's Joint Information Environment (JIE) and the IC's Information Technology Enterprise (ITE), respectively. These overview briefings provided a foundation for group discussion about adoption and prioritization of cyber resiliency techniques within these architectures.

The group next examined each of the cyber resiliency techniques in the context of how their use might affect the DoD JIE or IC ITE architectures, whether they increased cyber resiliency for the overall DoD or IC mission versus a local mission, and perceived adoption challenges. Finally, the group developed a set of high-level adoption principles and a notional prioritization of cyber resiliency techniques for the DoD and IC enterprise architectures, which Bryan Larish presented at the end-of-day workshop wrap-up session.

Applying Cyber Resiliency to Planned Architectures

Participants began the working session by defining a set of common characteristics of the JIE and ITE that the DoD and IC would need to consider when determining which cyber resiliency techniques would be most beneficial. Key characteristics included:

- **Culture change**: Both architectures reflect a shift from "stovepiped" systems to more enterprise capabilities.
- **Consolidation**: The DoD and IC could gain cost savings and other efficiencies by consolidating networks, applications, and data.
- **Enterprise security**: Both architectures reflect a relatively holistic approach to security that involves protecting networks and providing tighter/better managed access control.
- **Logical separation**: Separate data flows and systems within a single enterprise environment are common characteristics of both architectures.

With these common characteristics in mind, the group began to explore each of the cyber resiliency techniques in MITRE's Cyber Resiliency Engineering Framework, realizing that some of the techniques are more likely to apply than others and that their use must not interfere with the mission. The attendees agreed that the framework serves as a useful starting point for determining how and where to introduce cyber resiliency into future architectures such as JIE

and IC ITE. One suggested approach for determining which techniques show the most promise involves prioritizing mission-critical capabilities and threats to those capabilities and then prioritizing the application of cyber resiliency techniques accordingly.

Cyber Resiliency Techniques

The definitions of the techniques described in this section are taken from the Cyber Resiliency Engineering Framework developed by Deborah Bodeau and Richard Graubart of MITRE in September 2011.

Adaptive Response

> To practice Adaptive Response is to *take actions in response to indications that an attack is underway based on attack characteristics*. More specifically, Adaptive Response involves selecting, executing, and monitoring the effectiveness of the cyber courses of action that best change the attack surface, maintain critical capabilities, and restore functional capabilities.

Discussion centered on the recognition that "pulling the plug" in the face of a cyber attack is not an acceptable response. Decisions about how to react must take into account the local commander's role in a mission as well as the broader mission, and the different actions taken by each stakeholder should be well coordinated to achieve the most effective overall response. Roles and responsibilities for all involved must be well defined, with redundancies built in, and training plays an essential role. These activities must be well funded and repeated on a regular basis to keep pace with changing adversary tactics, techniques, and procedures and the addition of assets to the enterprise.

Analytic Monitoring

> To practice Analytic Monitoring is to *gather and analyze data on an ongoing basis and in a coordinated way to identify potential vulnerabilities, adversary activities, and damage*. To gather data, sensors are deployed within, and at the boundary of, distinctly managed sets of cyber resources. Coordination includes establishing coverage and timeframes or frequency for data gathering and analysis to avoid gaps or blind spots, and can include mechanisms for data fusion, correlation, and data mining. Examples of analysis include identifying anomalous behavior, performing malware analysis (passive, active, and postmortem), and using validation techniques to identify changes in infrastructure that indicate an ongoing attack.

The DoD/IC should use Analytic Monitoring to guide and increase the effectiveness of adaptive response.

Analytic Monitoring is most effective when systems gather the right kind of data at the right locations within the enterprise. This will require local collection of data as well as aggregation of

data for an enterprise. Attendees believed that the JIE provides monitoring for adversary activity; however, the DoD may have to place greater emphasis on the actions taken in response to detection of adversary activity. Command and control for response is not very well defined.

Coordinated Defense

> To practice Coordinated Defense is to *manage adaptively and in a coordinated way multiple, distinct mechanisms to defend critical resources against adversary activities*. Requiring the adversary to defeat multiple mechanisms makes it more difficult for the adversary to successfully attack critical resources, and increases the likelihood of adversary detection.

Coordinated Defense relies on a culture of shared situational awareness rather than on any particular technology.

Attendees agreed that both JIE and ITE would benefit from shared situational awareness. However, achieving such awareness requires a common understanding of the level of information to share and the partners with whom to share it.

Deception

> To practice Deception is to *use obfuscation and misdirection (e.g., disinformation) to confuse an adversary*. Deception can take the form of dissimulation ("hiding the real") or simulation ("showing the false").

Neither JIE nor ITE intentionally practices Deception today: techniques are still immature and costly to implement (over and above the security capabilities required). For these reasons, Deception as a resiliency technique is considered a low priority for both the DoD and IC architectures.

Diversity

> To practice Diversity is to *use a heterogeneous set of technologies (e.g., hardware, software, firmware, protocols) to minimize the impact of attacks and force adversaries to attack multiple different types of technologies.*

Attendees agreed that critical assets (platforms and applications) within JIE and ITE would benefit from Diversity. However, component standardization is ingrained; purchasing multiple different technologies and persuading organizations to maintain them would be difficult.

Dynamic Positioning

> To practice Dynamic Positioning is to *use distributed processing and dynamic relocation of critical assets and sensors*. Dynamic Positioning applied to critical assets will impede an adversary's ability to locate, eliminate, or corrupt

mission/business assets, and will cause the adversary to spend more time and effort to find the organization's critical assets.

Attendees believed that this technique could be relevant for the JIE Core Data Center, and would provide useful protections for "crown jewels" if used in combination with privilege restriction.

Dynamic Representation

To practice Dynamic Representation is to *construct and maintain dynamic representations of components, systems, services, mission dependencies, adversary activities, and effects of alternative cyber courses of action*.

This technique would enable live situational awareness of systems, infrastructure, and mission, but attendees concluded that it would be very difficult to build and scale, making it generally out of scope.

Non-persistence

To practice Non-persistence is to *retain information, services, and connectivity for a limited time, thereby reducing an adversary's opportunity to exploit vulnerabilities and establish a persistent foothold*. Non-persistence involves quickly refreshing information, services, and connectivity to known trusted states, and eliminating services, information, and connectivity that are no longer needed. Virtualization makes such refreshment much easier. Non-persistence is most appropriate when refresh is quick enough not to interfere with mission/business functions.

Both the DoD and IC make wide use of virtual machines, which could be an appropriate target for applying Non-persistence within JIE and ITE.

Privilege Restriction

To practice Privilege Restriction is to *restrict privileges required to use cyber resources, and privileges assigned to users and cyber entities, based on the type(s) and degree(s) of criticality and trust, respectively, to minimize the potential consequences of adversary activities.*

Both architectures already place a strong emphasis on this well-understood technique. JIE and IC ITE restrict privileges via runtime identity and access management (although it may not be applied to all cyber resources). Developing and managing policies for sharing data among organizations that need access represents one of the major challenges in both architectures.

Realignment

To practice Realignment is to *align cyber resources with core aspects of mission/business functions, thus reducing the attack surface.* Realignment

minimizes the chance that resources dedicated to activities that do not support mission/business functions could be used as an attack vector. One example of Realignment is offloading some less important cyber-supported functions to a service provider that is better able to support the functions, or eliminating certain data feeds or connections where the benefits of those feeds are determined to be less than the potential risks such connectivity imposes on the core mission/business functions.

Like Privilege Restriction, Realignment is a technique already familiar to JIE and ITE architects, who viewed it as standard Information Assurance engineering. Regardless, applying this technique would require effort to identify critical mission servers and ensure that they are not used for inappropriate or risky activities.

Redundancy

To practice Redundancy is to *maintain multiple protected instances of critical resources (information and services)*. These serve as backups in the case of localized damage to a resource and provide surge support when needed to support unexpected peak loads, faults, and failovers.

Attendees recognized the benefits of this technique in preventing single points of failure and as a strategy for avoiding or preventing cyber attacks. However, the DoD/IC would have to carefully balance applying Redundancy to systems, services, or data against a key consideration (for JIE in particular): consolidation. Additionally, infrastructure owners and owners of mission systems have different needs for Redundancy; thus, determining where best to apply this technique will present a challenge. In situations where Redundancy makes sense, attendees agreed that automated support for features such as backup and failover would be essential.

Segmentation

To practice Segmentation is to *separate (logically or physically) components based on pedigree and/or criticality, to limit the spread of or damage from successful exploits*. Segmentation reduces the attack surface and enables more cost-effective placement of defenses based on resource criticality.

Attendees saw value in the ability to segment "crown jewels" so that they could be isolated in the event of an attack. JIE in particular could provide Segmentation as a valuable service for specific users/missions.

Substantiated Integrity

To practice Substantiated Integrity is to *ascertain that critical services, information stores, information streams, and components have not been corrupted by an adversary*. Example mechanisms include use of integrity checks (e.g., checksums on critical records or software, use of the Trusted Platform

Module [TPM] to report system state information), data validation (checking that data conforms to its specified requirements, such as type or range), and tamper-evident technologies.

Attendees viewed the more traditional Substantiated Integrity techniques, such as Privilege Restriction and Realignment, as standard practices. Some of the more sophisticated methods for providing this capability (e.g., use of TPMs) would involve additional cost and infrastructure support, making it out of scope at this time.

Unpredictability

To practice Unpredictability is to *make changes frequently and randomly, not just in response to actions by the adversary*. Examples of unpredictable behavior include, but are not limited to, changing browsers and authentication mechanisms, encryption rekeying, and changing permitted ports.

As with some of the other techniques, participants viewed use of Unpredictability as costly to implement, but agreed it could be an important aspect of an overall defensive strategy.

Session Outcomes

Following the discussion of each of the cyber resiliency techniques, attendees developed five high-level principles for the DoD/IC to consider in adopting cyber resiliency techniques and a notional prioritization of the techniques to highlight those that would be most beneficial if applied now within the target architectures.

1. All of the techniques would be useful and beneficial, but some may be more difficult to adopt than others. For example, techniques such as Deception, Dynamic Representation, Realignment, and Unpredictability would be costly to implement and maintain.
2. Some of the techniques simply represent good engineering practice and are being incorporated into the JIE and ITE architectures. Both architectures already make heavy use of techniques that include Privilege Restriction and Substantiated Integrity (use of checksums and data validation, in particular).
3. The DoD/IC should consider cyber resiliency sooner rather than later. Some techniques, if not explicitly addressed early, will be executed poorly (see the Notional Prioritization below).
4. The enterprise approach embodied in the DoD and IC future architectures increases the importance of defining roles and responsibilities, especially in the deployment and management of cyber resiliency techniques. Incorporating cyber resiliency into these future architectures will require careful scoping and coordination between the mission and enterprise communities.
5. The DoD/IC must carefully consider potential impacts on the infrastructure and the mission before incorporating cyber resiliency techniques into JIE and IC ITE. They must

examine cost-benefit trade-offs to identify best fit techniques and calculate return on investment. Performing a cyber resiliency assessment is one way of determining which techniques would bring the most value to the enterprise architecture and where to apply them within the architecture.

With these principles in mind, and applying a degree of engineering judgment, attendees developed a notional prioritization of cyber resiliency techniques (see Figure 4). JIE/ITE enterprise architects would need to re-evaluate this list to confirm the suggested prioritization. Track participants placed techniques in a particular category on the basis of the perceived importance of the technique in improving the enterprise's cyber resiliency; they did not necessarily take into account other considerations such as cost and schedule impact. However, participants noted that the DoD/IC should take into account attacks from external actors as well as the threat from malicious insiders when they select techniques.

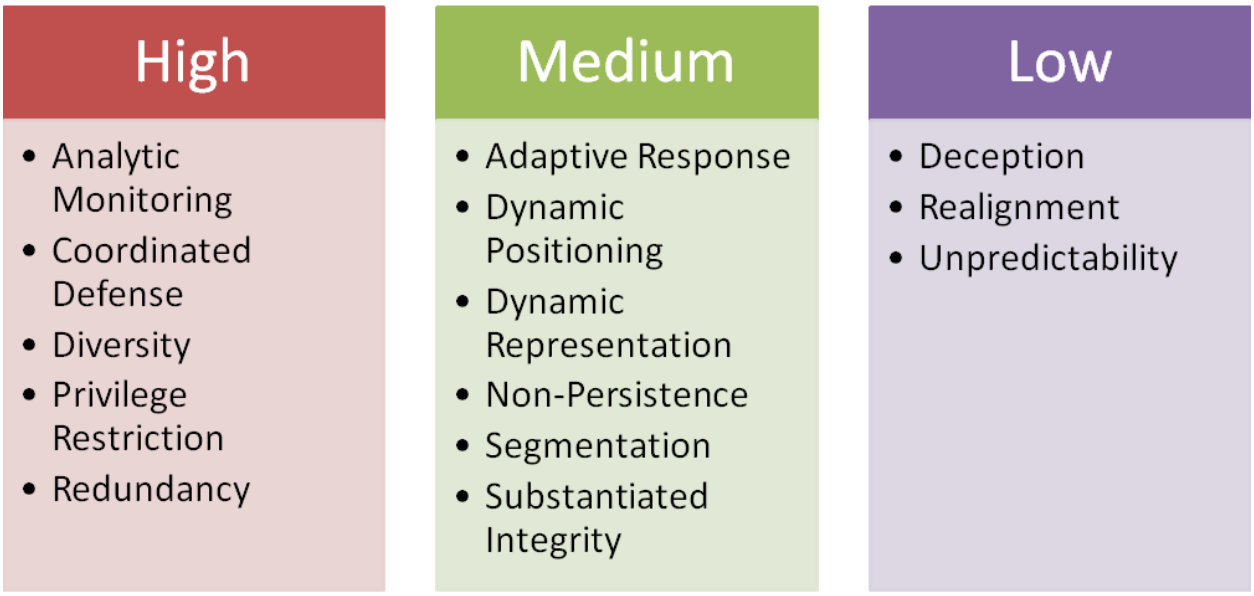

Figure 4. Prioritized Techniques

The DoD/IC can realize improvements in the overall cyber resiliency of its architectures through the judicious application of these technologies. The adoption principles and suggested prioritizations of cyber resiliency techniques should be brought forward to JIE and ITE Program Managers for consideration before the architectures are baselined.

Track 3: Advancing Cyber Resiliency Through Active and Adaptive Response

Track Chair: Steven Danko, NSA; sedanko@tycho.ncsc.mil
Co-chair: William Heinbockel, The MITRE Corporation; heinbockel@mitre.org

The third track focused on methods for applying static and dynamic response mechanisms that quantitatively improve cyber resiliency and measurably affect adversary behavior. The track began with a presentation about how NSA/R2 approaches the foundations of cyber resiliency. Afterwards, MITRE presented slides detailing a framework and vocabulary for response management. Then NSA and MITRE teamed to showcase the response management framework in the context of a cyber adversary vignette. The attendees then divided into teams to develop response vignettes focused on five cyber-related topics. At the conclusion of the track, each team described its results to all the track participants.

Foundations for Cyber Resilience

NSA launched the track by providing some motivation and framing of the problem, along with its perspective on a possible solution. That presentation noted that past practices in terms of cyber defense have not helped the DoD to better defend its networks. Even though we develop more and better sensors, we cannot win the battle against the APT. We believe that cyber resiliency can help us better defend our networks and "fight through" adversary attacks, but resiliency involves a more active function than simply detection: we need to know when, where, and how to act. We have garnered some initial capabilities and lessons learned in this area from previous research and discussions regarding "active defense."

NSA/R2 seeks help and feedback on the following:

- How can we identify and measure the effects of resilient responses on our users and missions?
- How might we choose to respond to attacks, and what knowledge is needed to support such a decision?
- What proportion of the decisions and responses should be automated versus human initiated (human in-the-loop) versus human oversight (human on-the-loop)?
- How can we define resiliency? How can we make it tractable and keep it independent from detection?
- How can we enable resiliency without having to identify the threat?

We already have some initial successes, in that our systems remain functional despite the presence of threats. However, while the systems continue to operate, decisions are made without awareness of the overarching mission and in an unknown, non-pristine state. We need further research to help in restoring systems while maintaining support to the mission, whether

by identifying and removing the compromises or dynamically restoring the system state from a known good configuration.

Significantly, our (defensive) actions are not at all resilient. Removing a system from the network does little harm to most adversaries and may have a tremendous impact on mission operations. Our response management process must also be resilient. Regardless of the actual responses, the response decisions must remain tactical. They must be chosen to best support the current and future mission needs, and we must be aware of the potential impact on both the threat and our own users. Furthermore, our actions should be unpredictable: adversaries should not be able to anticipate which response we will implement or to intentionally trigger a specific response.

Resilient Response Framework

MITRE presented a framework to assist in the discussion of resilient responses. Such responses represent a necessary evolution in cyber defense, as organizations initiate traditional "reactive responses" only after a lengthy investigation—which usually delays response until a few days after the initial incident. Additionally, since responses represent a reaction to a significant event, the response is of similar magnitude. This means that currently, cyber defense applies very powerful responses to all incidents; these responses usually harm operations more than the adversary.

Instead, MITRE proposes resilient responses: more frequent, more granular, and less invasive than the responses applied currently. However, responses alone cannot provide resiliency. Defenders need an overarching response management system to detect when to respond, select the most appropriate COA or COAs, and evaluate the result and side effects to inform future response decisions. A response management system monitors observables within the environment, aggregating metrics and response indicators. These metrics and indicators apply both before and after the response occurs. Response criteria include those metrics, indicators, or other measurements used to decide when to respond and determine which COAs are most applicable for the current situation. Effectiveness measures encompass the metrics, indicators, or other criteria used to determine how effective the response was and examine any side effects.

Table 2 lists the proposed working definitions for resilient response.

Table 2. Presented Working Definitions

Term	Working Definition
Cyber Resiliency or Resiliency	The ability of a nation, organization, mission, or business process to anticipate, withstand, recover from, and evolve to improve capabilities in the face of adverse conditions, stresses, or attacks on the supporting cyber resources it needs to function. [MTR 110237,

Term	Working Definition
	http://www.mitre.org/sites/default/files/pdf/11_4436.pdf]
Cyber Course of Action	A set of activities by cyber defenders and, as needed, other cyber staff and mission staff to address confirmed, suspected, or anticipated adverse conditions, stresses, or attacks on cyber resources. [MTR 110237, adapted] Note that performing no additional activities is a COA.
Resilient Response	A COA that provides or maintains cyber resiliency.
Criterion	A standard on which a judgment or decision may be based. [Merriam-Webster]
Response Criterion	A criterion for (i) determining whether or when to respond and/or (ii) selecting a response. A response criterion can be represented as a subset of the possible values of a metric, indicators, policy or practice criteria, or a combination thereof. Examples of response criteria include: • The number (number or count, an ordinal metric) of simultaneous users on an end system exceeds one (a threshold value). • The frequency (a ratio metric) of pings on a subnet doubles (a threshold value). • A remote file copy is followed by a remote execution command (an indicator). • A remote file copy is followed by more than one remote execution commands (a combination of an indicator and an ordinal metric). • An IAVA (Information Assurance Vulnerability Alert) is received from the Defense Information Systems Agency (a policy criterion for DoD organizations). • The time to implement a temporary patch exceeds the period before a planned software update that will include a permanent fix (a practice criterion, based on a response time metric). • The cost of the response (monetized cost, including lost productivity, offset by improved performance) is significantly less than the expected loss (again, monetized) from the adversary actions.
Effectiveness Measure	A measurement of the effects and effectiveness of a cyber COA. Effects can include impacts (positive and negative) on the mission or organization. Effectiveness can be expressed in terms of effects on the adversary and/or in terms of achievement of cyber resiliency objectives or sub-objectives.
Observable	A stateful property or measurable event pertinent to the operation of computers and networks. [STIX White Paper, http://stix.mitre.org/about/documents/STIX_Whitepaper_v1.0.pdf]
Cyber Threat Indicator Type	A pattern of observables (e.g., repeated Hypertext Transfer Protocol requests from the same Internet Protocol [IP] address).
Cyber Threat Indicator	A set of cyber observables combined with contextual information intended to represent artifacts and/or behaviors of interest within a cyber security context. [STIX White Paper]

Term	Working Definition
Measurement	A representation of a value, obtained by applying a system of measurement to the real world, *or* The process for obtaining a value. For the purposes of today's discussion, we focus on the representation of the value. As noted in Ford et al.,[1] measurements "range from weak to strong, with the weakest being nominal, and progressing though ordinal, interval, and ratio." Thus, a measurement can take the form of a metric (a quantified measurement, which could be ordinal, interval, or ratio) or an indicator (a nominal measurement).
System of Measurement	A system, consisting of a set or range of possible values and evaluation rules (including processes for obtaining data and algorithms for computing values), for deriving a value based on data gathered from the real world.
Metric	A quantified measurement, *or* The rule for obtaining the quantified measurement. For the purposes of today's discussion, we distinguish between a *metric value* and the *metric definition* (i.e., the range of values and/or the evaluation rule). For example, "The number of simultaneous users on an end system" is a metric definition; "3" is a metric value.

Working Sessions—"Future of Resilience"

The attendees were divided into teams to develop their own realistic cyber resiliency narratives or vignettes, focusing on one of five predefined areas: the "Internet of things," the mobile workforce, cyber-enabled transportation, safety-critical systems, and symbiotic systems. Each vignette consisted of a description of the particular system or scenario and insight into the challenge of maintaining system resiliency in the presence of persistent threats. Teams filled out a response vignette template (Figure 5) to describe the potential COAs that might improve system resiliency, such as the data necessary to determine when a response may be required, and to compare response options and options for evaluating response effectiveness.

[1] "Toward Metrics for Cyber Resilience," Proceedings of the 21st Annual EICAR Conference, http://nob.cs.ucdavis.edu/bishop/papers/2012-eicar/resilience.pdf

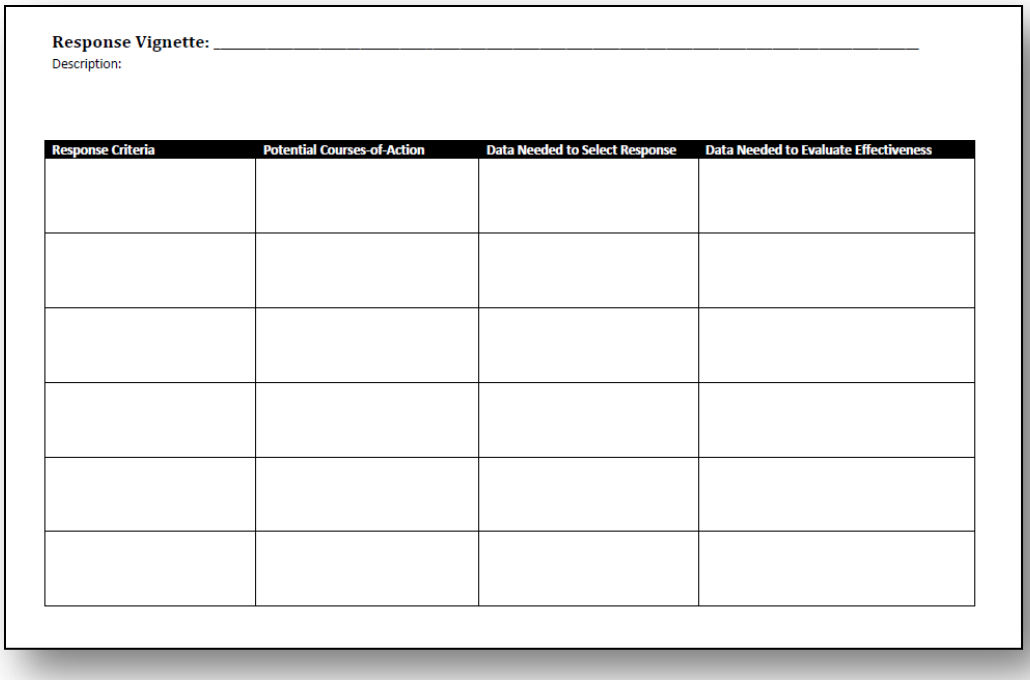

Figure 5. Response Matrix

Cyber Adversary

To illustrate the process, NSA and MITRE provided an initial vignette example that looked at a simplified timeline of a cyber adversary's life cycle from initial network exploit through foothold establishment, reconnaissance and expansion, to final detection and mitigation.

Vignette Summary

Description: An adversary has accessed your organization's network by sending a spear-phishing email with a malicious attachment, which was opened by a user. The adversary has just gained network access and is gathering reconnaissance on the environment and attempting to expand its foothold.

Table 3. Cyber Adversary Response Matrix

Response Criteria	Potential Courses of Action	Data Needed to Select Response	Data Needed to Evaluate Effectiveness
Execution of 2+ reconnaissance commands	• Call user • Deny execution • Isolate system • Increase data collection • Identify command shell point of origin	• DEFCON level • Cost • Time constraints • User history • Number and variety of commands executed • Confidence: presence of other detection criteria	• Command usage reduced • Subsequent observables • Impact on mission tempo, operations, users • Impact on threat: additional steps, shift to different system • User response—is there a legitimate explanation?
Local network recon followed by ping	• Drop ping; don't respond • Isolate source system • Increase data collection	• Capabilities present to drop ping/pong replies • System network history • Cost • Confidence: presence of other detection criteria	• Number of ping replies returned • Subsequent observables: – Does the adversary attempt to ping another system? – Does the adversary try another recon method?
Ping packets sent to multiple hosts	• Drop ping; don't respond • Isolate source system • Increase data collection	• Capabilities present to drop ping/pong replies • System network history • Cost • Confidence: presence of other detection criteria	• Number of ping replies returned • Number of systems sent ping requests • Subsequent observables: – Does the adversary attempt to contact another system? – Does the adversary try another recon method?
Samba (SMB) file transfer followed by remote execution attempt	• Call user • Disable SMB remote execution • Disable SMB • Delay execution • Deny execution • Migrate services • Increase data collection	• Executable file entropy, contents, properties, linked data link layers • User/host SMB history • System operational dependencies • Cost	• Status of executable • Time exec delayed (sec) • Subsequent observables: – Evidence of Trojan? – Further suspicious activity? – Check of why exec failed? – Number of re-attempts?

Response Criteria	Potential Courses of Action	Data Needed to Select Response	Data Needed to Evaluate Effectiveness
		• Confidence: presence of other detection criteria	• Subsequent SMB traffic • Impact on mission tempo, operations, users
Command shell originates at remote IP address	• Degrade QoS • Block remote IP • Deny execution • Call user	• Traffic associated with remote IP/subnet: traffic types, volume, entropy • Associated user history • Cost • Confidence: presence of other detection criteria	• Additional delay caused by degraded QoS (sec) • Number of packet time-outs or retransmissions • Time of command shell session (sec) • Traffic from remote IP/subnet: types, volume • Impact to mission tempo, operations, users
Execution of 1+ suspicious administrative commands	• Call user • Deny execution • Isolate system • Increase data collection • Identify command shell point of origin	• DEFCON level • Cost • Time constraints • User history • Number, variety of commands executed • Confidence: presence of other detection criteria	• User response • Number of additional commands executed • Subsequent observables • Impact on mission tempo, operations, users • Impact on threat: additional steps, shift to different system

"Internet of Things"

Topic: Many everyday objects are becoming more intelligent as a result of being equipped with processors, sensors, software, and the ability to communicate. Eventually, these objects, or things, will lead to a massive expansion of the Internet and transform daily life in such areas as continuous medical care, emergency response, home automation, and public utilities. While these new things will make life more convenient, they will also elevate the risk of sabotage and malfunction. What incremental responses might increase the system resiliency of the Internet of Things?

Vignette Summary

Description: A terrorist attack has used Internet-enabled devices to disrupt or overload the power grid. The grid not only provides convenience in terms of lights, cooling/refrigeration, and so on, but also powers critical systems in hospitals and elsewhere (usually protected through redundancy or backups). Furthermore, power comes from multiple sources (e.g., coal/nuclear

power plants, wind, water) and is shared via interconnected grids to create resiliency over entire geographic regions.

Table 4. "Internet of Things" Response Matrix

Response Criteria	Potential Courses of Action	Data Needed to Select Response	Data Needed to Evaluate Effectiveness
Increase in customer complaints	• Increase monitoring	• Complaint specifics • Known bugs • Supply chain details • Complaint frequency	• Complaint frequency
Inconsistent physical indicators from individual device(s)	• Fix • Alert manufacturer • Contact customer	• Intelligence from local device • Environment sensor data • Cross-device correlation	• Trending of number, severity of problems • Trending of number, severity of affected devices
Change in power usage (regional)	• Shed some load • Shift power load • Contact owner	• Power usage data partitioned per device • Baseline of power usage • Bandwidth trends	• Percentage of power load change
Activity on deployed decoy	• Delay actions on decoy • Warn • Redirect other traffic	• Normal activity (on non-decoys) • Activity specifics • Frequency • Severity	• Time spent • Activity trends on decoy
No power	• Fix (draw power from other sources) • Notify • Reconstitute • Evaluate area	• Power level • Power loads of other potential power sources	• Time to fix • Change in power levels

Mobile Workforce

Topic: More and more enterprise computer systems are now mobile and wireless, which makes them difficult to secure. At the same time, organizations have become more flexible, creating dynamic work environments where employees work from almost anywhere and even bring their own mobile devices to work. As we increasingly use the same devices at home and at work, how can resilient response reduce the expanded risk posed to private data of organizations?

Vignette Summary

Description: A large regional disaster (e.g., earthquake, tsunami) has occurred. International aid and assistance arrives to help treat the wounded and restore services. Terrorists attempt to manipulate cell towers and user devices to undermine the aid efforts by inserting misinformation to influence planning. A resilient system should have the ability to prioritize the mission (rescue efforts) and the increased load, and attempts to mislead the system should not achieve any sustained effect.

Table 5. Mobile Workforce Response Matrix

Response Criteria	Potential Courses of Action	Data Needed to Select Response	Data Needed to Evaluate Effectiveness
Increase in the number of failed cellular connections **Cellular configuration changes** **Multiple cell tower activity** **Large, unexpected changes in device locations**	• Lock/block mobile account(s) • Contact configuration team • Proxy communications and increase monitoring • Degrade confidence • Contact vendor(s) • Validate data against that from alternative sources • Deploy redundant infrastructure	• Network traffic levels, trends • Login information • Provider data activity • Mission information (system and data criticality) • Availability and trustworthiness of alternative data sources	• Number of people/area affected • Time to respond, restore network • Accuracy of alternative data sources

Cyber-Enabled Transportation

Topic: Trains, planes, and automobiles are becoming Internet enabled and increasingly autonomous. Any disruption may have costly consequences. How can we prepare cyber-enabled modes of transportation to withstand the potential threats?

Vignette Summary

Description: While traffic control systems are currently associated with airplanes, cyber-enabled cars will create the potential for autonomous, inter-car traffic control. In this case, an attacker may be able to use low-and-slow attacks to take control covertly. In such cases, defenders need a long data history for correlation to protect the vehicles and passengers from harm. Additionally, any detection of factors related to resiliency must accommodate natural faults and environmental factors (e.g., slippery roads, inclement weather, flat tire). Cyber-enabled transportation must be treated as a safety-critical system; resilient systems must enable fallbacks to a predefined “safe” state.

Table 6. Cyber-Enabled Transportation Response Matrix

Response Criteria	Potential Courses of Action	Data Needed to Select Response	Data Needed to Evaluate Effectiveness
Threshold outside of or at tolerance limit	• Increase data collection • Perform root cause analysis • Contact manufacturer for known issues, debug problem systems • Disconnect cyber components • Fall back to a manual, safe mode (e.g., disable cruise control)	• Criticality of system (mission, passengers, count) • Type of failure (accidental or intentional) • Type of data and thresholds (static or dynamic) • Environmental conditions (weather, traffic, geolocation)	• Did the tolerance decrease? • Rate of threshold change
Continued/ persistent threshold violations	• Isolate system • Recall or ground transport • Modify algorithms • Attempt to attribute (sensor failure?) • Enable self modification • Revert to external control • Change route • Correlate with other, similar systems	• Duration of violations • Presence of control modules • Presence of redundant or diverse devices (global positioning system [GPS] vs. compass)	• Rate of violations • Rate of threshold change • Number of similar systems affected

Safety-Critical Systems

Topic: Safety-critical (or life-critical) systems are those whose failure may result in death or severe damage. The design of such systems must ensure that if they fail they can still continue to operate (e.g., elevators) or become safe (medical devices), secure, or passive (aircraft landing system). What capabilities can defenders incorporate into safety-critical systems that will increase system resiliency without compromising system reliability?

Vignette Summary

Description: Healthcare abounds in safety-critical systems. For example, failure of a pacemaker could result in the death of the patient. Many of today's pacemakers incorporate remote sensing and configuration capabilities (RCSs) that doctors can use to collect and monitor the devices and make any necessary adjustments to their configuration. These devices are already known to be susceptible to microwave and magnetic influences; they can also be intentionally threatened by adversaries.

Table 7. Safety-Critical Systems Response Matrix

Response Criteria	Potential Courses of Action	Data Needed to Select Response	Data Needed to Evaluate Effectiveness
Control system receives indication from pacemaker that condition exists that requires remote response	• RCS triggers further request for diagnostic reports (always?) • Call patient to order him/her to hospital • Call patient to get further data • Send paramedics, call 911 • No action	• Diagnostic information • Correlation data for patient history (indicating either device or patient anomaly or not) • Correlation data across patients with same device • Authentication information • Location information	• 911 intervention or hospital visit saves life vs. unnecessary • Inconsistency of information provided (e.g., diagnostic data vs. alert; authentication failure; location data vs. actual location; historic vs. event)
Failure to receive report from device to RCS when expected	• Always: RCS pings device (how many times?) on first "miss" or after a threshold of "misses" • Out-of-band contact with patient • Policy: patient pre-notifies that s/he will be out of range at certain times • Change communication channel • Contact manufacturer	• Redundant location information (e.g., GPS data) over time • Historical data on communication errors	• Was communication restored? • Was device behavior restored? • Patient mortality, mortality rate • Time to restoration

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013
Approved for Public Release. Case Number 13-4210. Distribution Unlimited.

Response Criteria	Potential Courses of Action	Data Needed to Select Response	Data Needed to Evaluate Effectiveness
	• Contact emergency services/police to investigate		
Device receives disruptive/risky/ harmful commands (e.g., to harm patient)	• RCS detects anomalous command (sequence) • Device detects anomalous command (sequence) • RCS detects device response to command (sequence) it did not send	• History of device status reports (correlation with RCS issued commands) • Profile of safe change patterns/ thresholds against which to compare actual change • History of normal/ predictable command sequences • Authentication information	• Discovery of malicious command • Ability to re-establish proper behavior • Patient physical response (up or down)
Device behaves normally but patient is in trouble	• Out-of-band sanity protocol • Fault injection on periodic or irregular basis		
Device fails to respond to RCS commands			
Device fails to receive regular ping from RCS			
Device reports unexpected sequence of behavior change			

Symbiotic Systems

Topic: The overall resiliency of a system is not necessarily the sum of the resiliency of its individual components. Some components may have interdependencies with other components that an adversary could manipulate to compromise the overall resiliency of the system, regardless of the resiliency of each component. In such cases, how can incremental responses help to ensure symbiotic resiliency?

Vignette Summary

Description: One example of a symbiotic system is a home security system. Such a system has dependencies on many external factors, including the power grid, communication networks (Internet, telephony, or cellular), the monitoring company and its ability to detect and respond, the user/homeowner, and emergency response services. This wealth of dependencies means that these systems are usually built with multiple, diverse paths such as multiple power and communication options. Additionally, a web of potential communication may exist among the security system, homeowner, monitoring company, and emergency responders. A break or delay in any one of these must not cause failure of the entire system.

Table 8. Symbiotic Systems Response Matrix

Response Criteria	Potential Courses of Action	Data Needed to Select Response	Data Needed to Evaluate Effectiveness
Power failure	• Contact the power company • Switch to backup (battery, uninterruptible power supply) • Notify user • Notify monitoring company	• Availability of redundant or diverse power options • Battery power level • Duration of power outage • Area affected (individual house vs. immediate area) • Current weather and environmental conditions • Any nearby repair work	• Frequency of outages • Duration of affected operation • Time needed to isolate cause • Time taken to restore primary power source • Time taken to notify user, monitoring company
Audible alarm	• Neighbors respond • Contact homeowner • Contact emergency response	• Home vicinity (any nearby neighbors) • Homeowner availability • History of neighborhood crime and response	• Number of reports from neighbors regarding alarm • Time to contact the homeowner • Time to contact emergency response
Disabled alarm	• Notify user • Require user verification • Require out-of-band user verification	• System usage patterns • History	• Number of verification attempts • Number of out-of-band verification attempts • Time taken to verify user
Sensor failure	• Notify homeowner • Notify system	• Sensor, battery age • System history	• Time to restore sensor • Time to notify user

This report was compiled from sessions held at the MITRE Third Annual Secure and Resilient Cyber Architectures Workshop, 2013

Response Criteria	Potential Courses of Action	Data Needed to Select Response	Data Needed to Evaluate Effectiveness
	technician • Trigger alarm • Monitor utilities	• Local events (weather, crime) • System usage patterns	
System cannot connect to monitoring company	• Alert homeowner • Fallback to telephony, cellular, or backup communications • Contact emergency response	• System history • Local events (weather, crime) • Urgency • Active alerts • Phone line availability • Cellular signal strength	• Time to restore

Notable Findings

As the Internet becomes more prevalent and mature, our physical and cyber worlds become more intertwined. While more "smart" and cyber-enabled products may simplify our lives, these products confront a new, global threat emanating from the Internet. Cyber resiliency is as critical as ever, and defenders are just beginning to learn what resiliency techniques can achieve.

One theme common to all the vignettes was time. The ability to observe an environment and monitor changes and relative rates of change was an important criterion. The teams often measured response effectiveness in terms of the time to react/notice the issue or the duration of the issue before resolution.

All the teams also noted the dependence of resiliency on information and infrastructure. The richer the environment, the greater the potential for resiliency. However, defenders' need for data to analyze and correlate increases the importance of data ownership, availability, and usage. Once defenders know what data they need in order to determine response criteria or to evaluate response effectiveness, they must actually obtain that data. This should be trivial within their own infrastructure, but more complicated with external infrastructures. Enhancing resiliency will require collaboration across many organizations and will place a high value on establishing trust, authorization, and reliable cross-domain data sharing.

The participants also noted several other interesting possibilities for approaching resiliency and enabling COAs:

1. *Must resiliency always be built in? Are there technologies or environments that have inherent resiliency—COAs that do not have to be built into the system architecture?* For example, a critical mass of cyber-enabled cars would form a "smart car grid." If one driver loses control, the surrounding cars can detect this and take over. Built-in COAs are largely a precautionary measure that may get little use. If cyber defenders can identify and leverage the inherent resiliency, they can pursue the most cost-effective solutions first.
2. *How can we extend the concepts of risk management into the resiliency realm?* Risk can be introduced into a system under controlled conditions to better evaluate resiliency under certain conditions. However, in some circumstances, this may expose the system to additional risk. Defenders can introduce external stressors for safety-critical systems such as pacemakers to establish actual baselines and evaluate the resiliency of the system once it has been implanted in a patient.
3. *How can we actually measure resiliency?* Resiliency is highly contextual. The value and reliability of thresholds and measures can vary drastically from one environment to the next. Certain response criteria or effectiveness measures may only apply in certain situations. Therefore, how can we collaboratively research, develop, and share resiliency and response management techniques in such a way as to minimize the overall level of effort?
4. *Are there response criteria and effectiveness measures that do not involve time?* While the element of time was a prevalent theme, it requires substantial data collection and correlation to generate criteria. Defenders can use historical information for baselining, but this demands a collection of all data for the period, and the data should be refined to be most appropriate for the usage context. Calculating duration is easier, but still requires tracking the start and stop events to determine the elapsed time. This, in turn, highlights the problem of defining "start" and "stop" for cyber events. Should defenders look back to the first detection of the adversary, determine or guess when the adversary first gained access, or start with the first time adversary actions had an impact on the system? While these decisions may seem trivial, events must be measured consistently and according to the same criteria. Identifying criteria that involve less complex analysis will help to simplify this task, reduce the overhead, and improve responses and overall resiliency. Defenders should prefer less complicated criteria and measures, such as single indicators or simple event sequences, to time-based ones.

Value Proposition

The Cyber Resiliency Workshop continues to be valued and to serve the community well. This is evidenced by both the increase in attendance over the last three years and the fact that the workshop committee has already received a number of requests and queries about next year's workshop.[2]

This year's workshop demonstrates the relevance and cross-disciplinary applicability of cyber resiliency and underscores the perceived value of incorporating cyber resiliency techniques into new and existing architectures and environments. In previous years we saw resiliency applied within the commercial sector; this year we learned that there is perceived benefit in applying cyber resiliency to JIE/IC ITE and for use in critical infrastructure protection. This year's workshop provided members from these diverse communities an opportunity to discuss and begin to reconcile their different cyber resiliency needs and concerns, and to identify new ways to address adoption barriers (i.e., economic and policy).

In addition, the workshop continues to be a vehicle for identifying and setting direction for addressing cyber resiliency needs. At last year's workshop, MITRE introduced its Cyber Resiliency Engineering Framework. During this year's workshop, all of the talks and tracks featured discussions based on the framework, resulting in a set of outcomes and recommendations intended to move the community forward in their actions and thinking relative to cyber resiliency and how it can be applied effectively. Examples include an activity to review the security controls in NIST SP 800-53 to understand the degree to which they address cyber resiliency, and a notional prioritization of cyber resiliency techniques for application within the JIE and IC ITE architectures.

In conclusion, the Cyber Resiliency Workshop planning team has high expectations that the fourth annual workshop will continue to provide value to the community and anticipates that it will take place in the first part of 2014.

[2] Individuals and organizations that are interested in participating in the 2014 workshop are encouraged to contact the Cyber Resiliency Workshop Planning Team: secureandresilient@mitre.org.