2nd Annual Secure and Resilient Cyber Architectures Workshop

Final Report

May 31 – June 1, 2012

Table of Contents

Topic Summaries

2nd Annual Secure and Resilient Cyber Architectures Workshop

The following records the results of the discussions and follow-on interactions of the subject tracks from the 2nd Annual Secure and Resilient Cyber Architectures Workshop held on May 31-June 1, 2012 at the MITRE facilities in McLean, VA. These summaries, one from each of the three tracks – Framework, Metrics and Assessment, and Engineering Principles – were reviewed by attendees from each of the respective tracks through email exchanges. They summarize what the track attendees believe to be the most salient parts of the discussion, items of consensus, questions posed from the discussions, and comments on next steps.
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

Comments from readers are welcome through the contact email address:
secureandresilient@mitre.org

The Cyber Resiliency Workshop Committee
21 September 2012

Cyber Resiliency Framework & Terminology

Track Chair: Ron Ross, NIST, ron.ross@nist.gov
Co-chair: Richard Graubart, The MITRE Corporation, rdg@mitre.org

Background

Underlying the Framework track was the understanding that there are dozens of definitions of resiliency and cyber resiliency available. Rather than try to anoint one of these or create another, this track was focused on trying to establish a (relative) consensus on the nature of a framework and supporting terminology for cyber resiliency. Potentially this may set future direction of cyber resiliency guidelines and provide an ongoing forum for discussion of cyber resiliency issues.

Scope of Cyber Resiliency

Track participants looked at the question of what is the scope of cyber resiliency and considered various options:

1. Ensuring resiliency of the cyber elements on which a mission or business function depends from attacks only of a cyber-nature;
2. Ensuring resiliency of the cyber elements on which a mission or business function depends regardless of the nature of the attack; or
3. Ensuring resiliency of non-cyber elements from cyber based attacks.

Of the three above, we seemed to agree that the first was too narrow. We found that the second and third were both in scope. Thus cyber-attacks on non-cyber elements are in scope as are non-cyber-attacks[1] on cyber elements. Only non-cyber-attacks on non-cyber elements are out of scope.

In discussing cyber resiliency, we considered various timeframes:

1. Resiliency after an adversary successfully breaches perimeter defenses and then compromises critical missions and/or business functions;

[1] However, the emphasis is on cyber-oriented attacks.

2. Resiliency after an adversary successfully breaches perimeter defenses; and
3. Resiliency to preclude an adversary from breaching perimeter defenses.

The consensus appears to be advocating for a full spectrum view of the cyber kill chain. There's a danger in stovepiping these concerns, so we need to be concerned about resiliency at all stages in the cyber kill chain: after mission/business compromise; after a breach of the perimeter; and before a breach of the perimeter. Resiliency includes trying to keep the adversaries out, but it must not stop there. It has to also address the need to respond if and when adversaries are successful in breaching the perimeter.

Changes in Mindset

The community of thought leaders, policy makers, architects, engineers, and researchers committed to building more secure and resilient architectures have to promote a change in mindset. Therefore, all stakeholders must recognize that:

1. Even with correct implementation of all the necessary perimeter-based security, and all the continuous monitoring to ensure that patches are applied and vulnerabilities are closed, the adversaries may still breach the IT infrastructure; and
2. Organizations must plan for, and be able to execute a mission-continuity response to attacks after perimeter defenses have failed.

But there is a balance we have to maintain. If we focus on actions to take after the perimeters fail, we may cause people to think that perimeters don't matter. We need to expand traditional information security with resiliency (i.e., boundary defense is a minimum necessary requirement but it is not sufficient to protect organizational missions and business functions). Moreover with architectures now including concepts such as mobile computing and cloud, and supply chain and insider attacks, the concept of securing a perimeter becomes less meaningful, and the need to "harden" mobile systems (consistent with their relative mission criticality) increases.

We may be able to promote this change in mindset by use of a healthcare arena analogy. People can exercise daily, eat balanced meals, and wash their hands. But they still get infections, cancer, and heart attacks. That does not mean one should not do these proactive measures (such as hand washing[2]). Nor does it mean those measures are not beneficial. It only means that they not sufficient. Hence, healthcare invested not only in preventative medications and techniques, but also in reactive medications and technique such as anti-cancer research/medication, post-heart-attack medication and treatment, etc. Cyber resiliency similarly needs to expand the scope of security beyond the proactive means of keeping an adversary out, but not at the expense of perimeter defenses.

[2] Follow on point: it was noted that even among healthcare professionals one cannot achieve anywhere near 100% hand washing compliance (number mentioned was 80%). The point to be taken is that it is not realistic to assume that an organization's proactive and preventative security measures will always be carried out.

Where Does Cyber Resiliency Fit In?

The track members talked about the relationship of cyber resiliency to other disciplines; in particular, resiliency engineering, cyber security, and mission assurance. Several good discussions and agreement arose from this. One point involved one of the subsequent visuals which showed cyber resiliency as sitting in an overlapping Venn diagram with other disciplines. The track members came to the agreement that was incorrect. Rather, cyber resiliency is a bridge filling a gap between these three disciplines (see Figures 1 and 2).

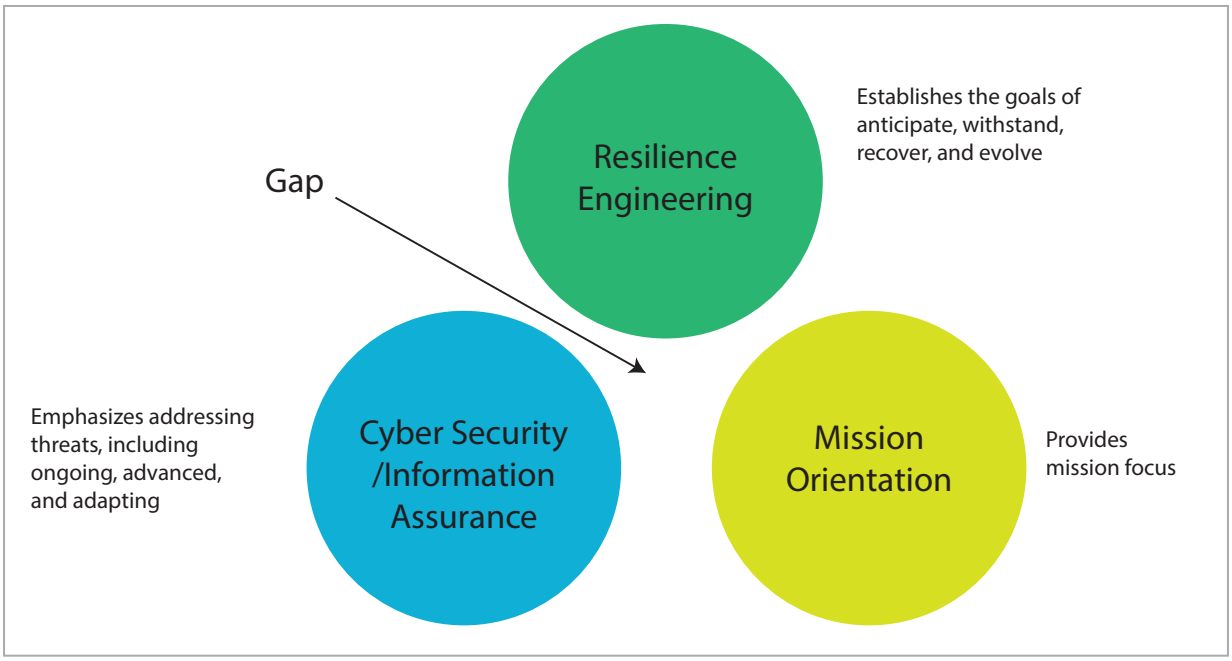

Figure 1

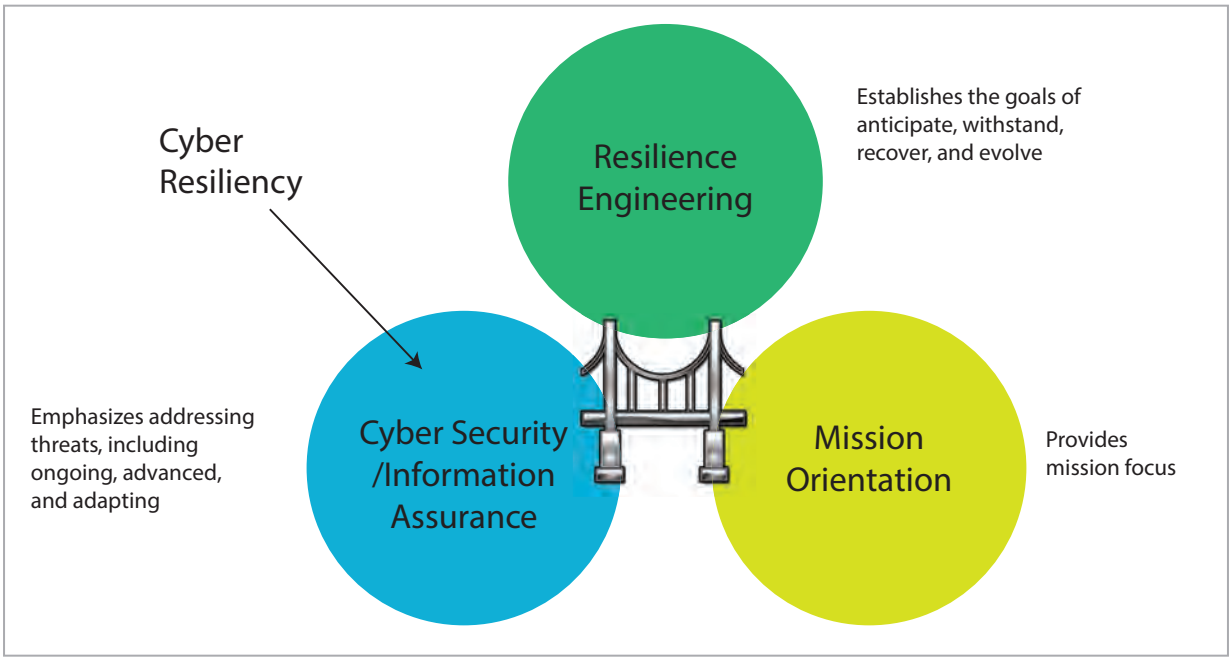

Figure 2

A second point was that cyber resiliency against cyber threats is different than resiliency against natural disasters. Natural disasters are one-time events, or in the case of earthquakes predictable follow-on events (aftershocks, tsunamis, etc.). Moreover, natural disasters do not target specific structures or entities (although some, especially if not well constructed, may be more vulnerable to such events). Natural disasters do not evolve or change in response to activities by defenders to mitigate the event. The need to deal with multi-pronged, targeted attacks that evolve and change in response to defender actions makes cyber resiliency different from the established resilience engineering discipline.

A third point was on the need for a mission focus. In the past, some have approached cyber security with a system focus and ignored the larger mission context. Sometimes a system must be sacrificed to preserve the mission. To be effective cyber resiliency needs to be done with a mission focus in order to ensure mission resilience to cyber-attacks. Moreover, focusing on the mission helps prioritize the possible resiliency measures so that effectiveness (e.g., in terms of mission resiliency) can be balanced as well as possible with efficiency (e.g., in terms of affordability).

A fourth point was that the cyber resiliency diagram MITRE presented was engineering focused. Engineering is certainly a key factor, but there are other factors to be considered to achieve cyber resilience. These include architecture, operations, and acquisition (see Figure 3 below). All of these factors need to be considered over the entire life cycle. These factors do not work in isolation and there is a feedback loop among the factors.

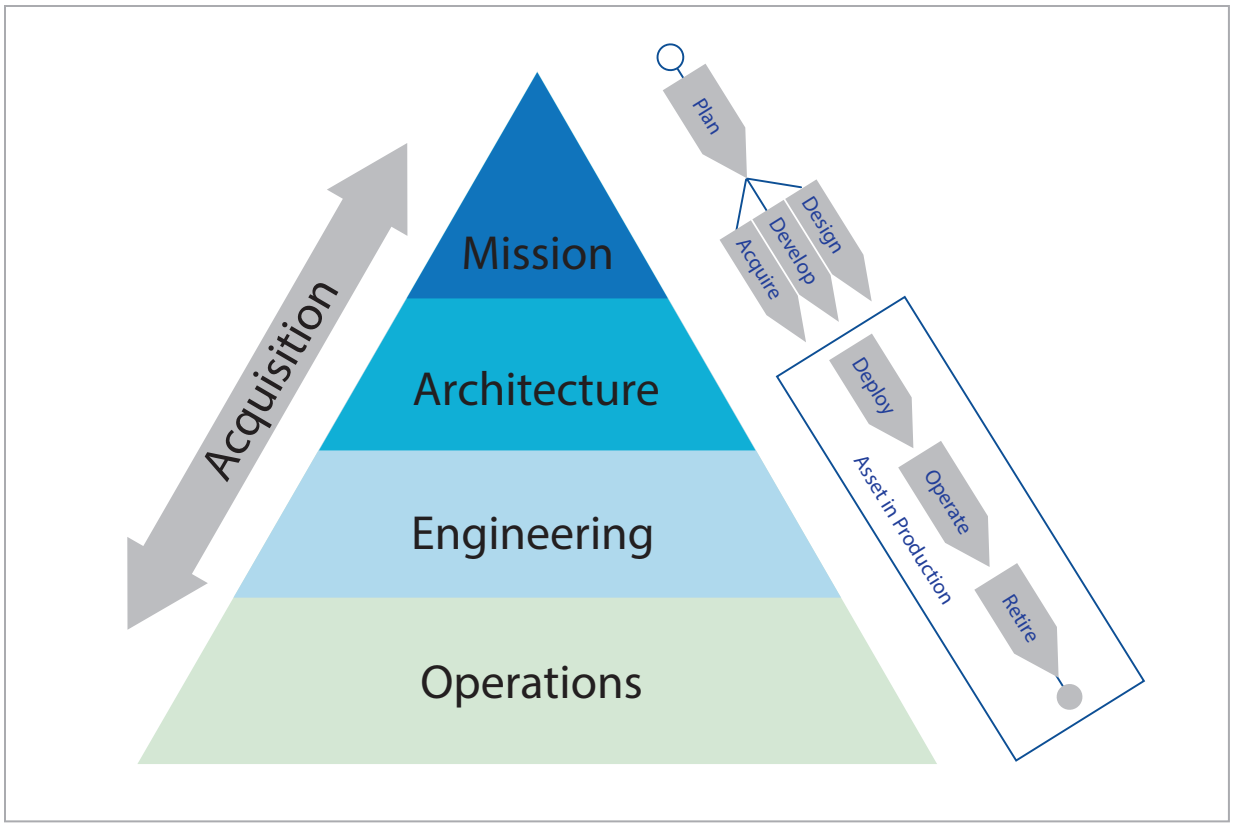

Figure 3[3]

[3] Right-hand side of figure 3 provided by SEI.

Examinations and Interactions of Various Frameworks

As part of the track, we looked at various frameworks: MITRE's Cyber Resiliency Framework, Software Engineering Institute's (SEI)'s Resiliency Management Model (RMM), and Johns Hopkins University/Applied Physics Lab (JHU/APL) Mission Based Analysis (MBA) methodology. With regard to the MITRE framework, there appeared to be agreement that the proposed goals, objectives, and practices were reasonable. That is not to say it was considered complete or correct, or that participants agreed on all aspects of the taxonomy (e.g., some questions arose as to whether dynamic positioning should be its own practice or part of deception).

We also looked at the resiliency model from CERT-RMM. The model uses the same architecture as CMMI. Operational resilience is the emergent property of an organization that can continue to carry out its mission in the presence of operational stress for a specific period of time. The organizational context includes productive activities with assets in production. The four types of assets are people, information, technology, and facilities. Resilience needs to be built at the assets level, a clam shell that covers both sides of the asset. Resilience management covers the life cycle of an asset.

CERT-RMM has 26 process areas in four categories. They are the enterprise foundation for these activities. This is a holistic, broader view. In this model, there are 251 practices spread across 26 process areas.

We also looked at the JHU/APL Mission-Based Analysis methodology and saw how it interacts with the MITRE Cyber Resiliency Framework.

Follow on work is needed to further examine relationships between the various frameworks and methodologies. But the initial view is that there is considerable synergy between them (see Figure 4 below).

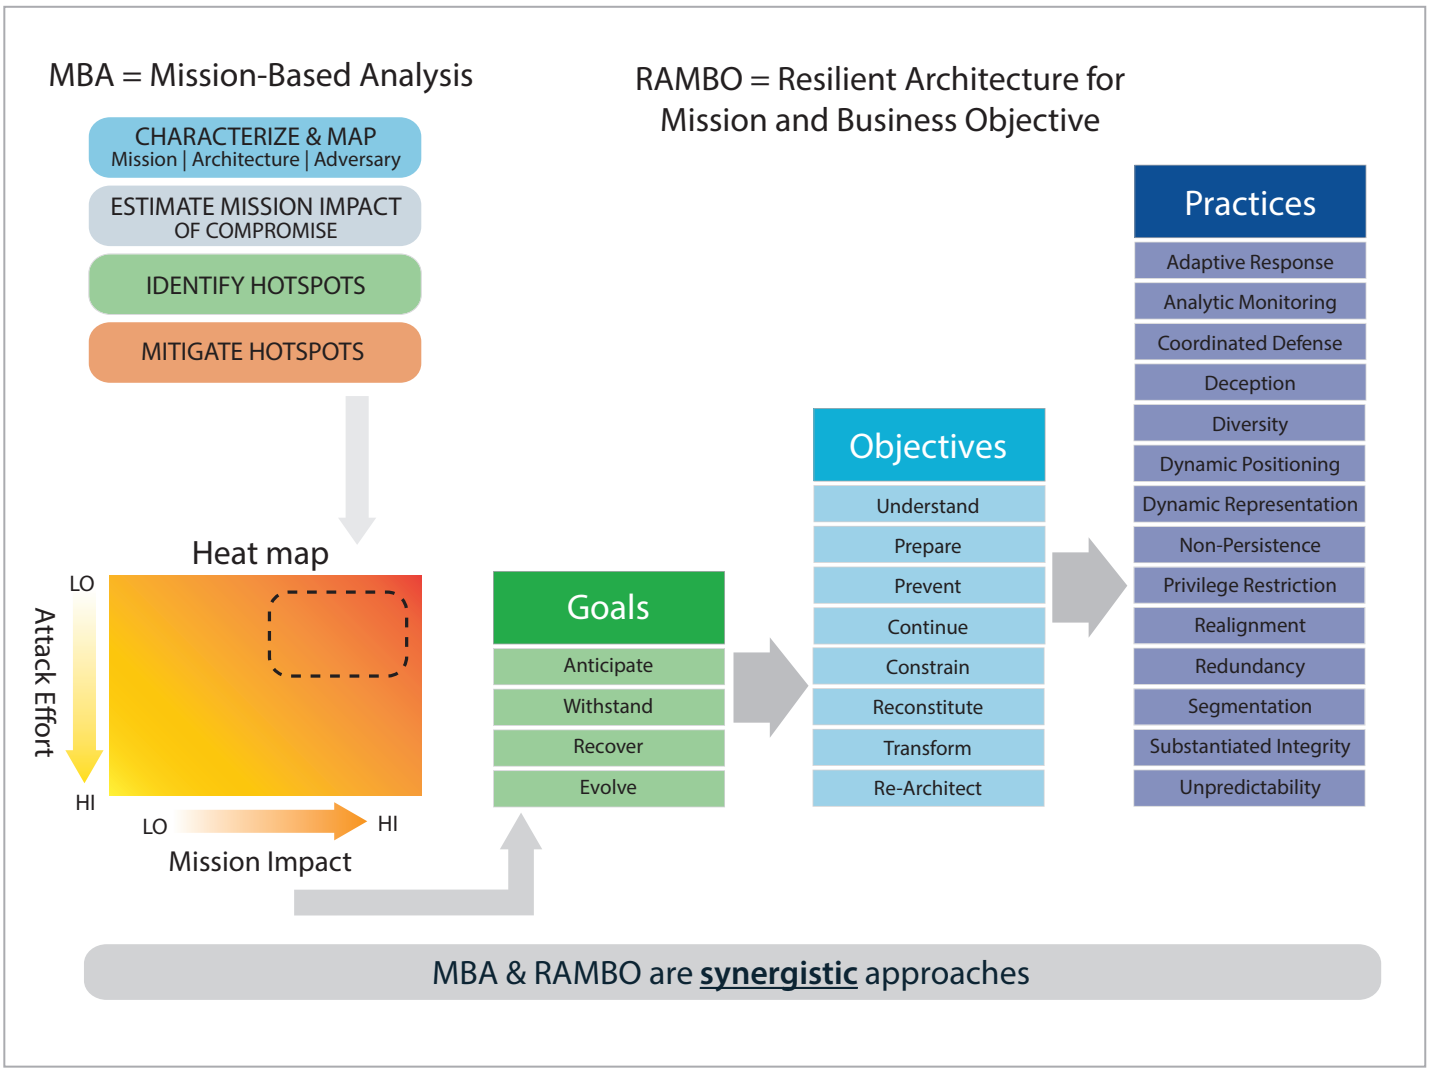

Figure 4[4]

Next Steps

Develop a unified, mission-driven, resiliency framework (architecture, engineering, operations, and acquisition):

- Integrate goals, objectives and practices from various sources into the framework;
- Allocate/map cyber resiliency practices across the framework;
- Develop NIST SP800-53 security controls related to resiliency practices;
- Define a strategy process to deploy controls into all layers (mission, architecture engineering and operations[5]); and
- Develop investment decision support tools and methods, and share data (cost, schedule and performance).

[4] Note: In Figure 4 what is referred to as RAMBO is what is elsewhere referred to as the MITRE Cyber Resiliency Framework.

[5] Part of the activity would also involve revisiting the set of layers to ensure that all are represented. Another possible layer would be one reflecting the human element.

Metrics and Assessment for Secure and Resilient Cyber Architectures

Track Chair: Nadya Bartol, Booz|Allen|Hamilton, bartol.nadya@bah.com
Co-chair: Deb Bodeau, The MITRE Corporation, dbodeau@mitre.org

Background

The announcement for this track read: *Something needs to be measured, or there will be no basis for telling whether investments or operational decisions are having any effect. Different approaches can be taken to defining resiliency metrics – e.g., Goal-Question-Metric, defining metrics based on temporal models (pre-event, trans-event, post-event). Metrics can be defined for capabilities and behaviors at different levels – organization, mission or business process, task or mission function, system, technology. Metrics can serve the needs of different stakeholders, from cyber defenders to product vendors to mission or business owners. We will focus on practices for defining, evaluating, and tracking resiliency metrics, and for sharing resiliency-related data, metrics, and lessons-learned. The track will identify the top principles and lessons-learned, to improve the state of practice.*

The track included an initial discussion, focused on how general principles of assessment and metrics apply to the cyber resiliency domain; presentations by participants on resources for resiliency assessment and metrics; identification of questions to motivate the development of metrics, based on the use case-driven analytic approach presented; and development of top principles and hard problems.

Challenges

Participants identified the close relationship among the three tracks – Framework and Terminology, Metrics and Assessment, and Engineering Principles – as an impediment to being able to make substantive progress on what to measure for cyber resiliency. To develop appropriate goals-based metrics participants would need the results from discussions at the Framework and Terminology and Engineering Principles tracks. These results would create goals and objectives for the Metrics track to work with to create metrics. However, the participants agreed that progress can be made in discussing the general principles for resiliency measurement which can later be used to apply to the results from the other two tracks.

Participants noted the community has not yet achieved consensus on a single definition of resilience (whether modified by "cyber" or not). However, agreement on a definition is not a prerequisite to defining metrics – the discussion of principles for how to measure resilience and what kinds of measures may be meaningful will help move the general discussion about cyber resiliency engineering forward. As was noted in the presentations on the first day of the workshop, the relationship between cyber resiliency and security[6] is still being articulated. In some ways, talking about cyber resiliency forces the community of practitioners to return to the fundamental objectives of security: confidentiality, integrity, availability, and accountability (in the sense of being able to present a realistic account of what has occurred). However, the concepts of continuity and recovery are central to resiliency, so that integrity and availability take higher importance than in the security problem domain, where confidentiality has historically taken precedence. Resiliency means continuing to operate under attack and recovering to a known state with known losses. In the context of cyber resiliency, measurement is essential to determining what can be trusted, as a basis for recovery.

Key Inputs

Key inputs to the discussion included

- A presentation by Nadya Bartol on a general process for creating and maintaining a security metrics program. The presentation referenced NIST SP 800-55 Rev1 (http://csrc.nist.gov/publications/nistpubs/800-55-Rev1/SP800-55-rev1.pdf), ISO/IEC 27004, ISO/IEC 15939, CMMI and CMMI Goal-Question-Indicator-Measure (GQIM) and the PracticalMeasurement Framework for Software Assurance and Information Security (https://buildsecurityin.us-cert.gov/swa/downloads/SwA_Measurement.pdf.) Participants concurred that this general process is relevant to cyber resiliency as well as security. Discussion noted the importance of feedback, so that the set of metrics used can be evolved based on experience of creating and using metrics as well as evolving organizational needs, goals, and environment. Tools and processes are needed to ensure that the reliable data needed to support construction of the metrics are available. Thus, metrics and the metrics process must change to adapt to the evolving stakeholder needs and to facilitate learning from metrics.

- A presentation by Julia Allen on Measuring Operational Resilience. Extensive materials on the CERT Resilience Management Model (CERT-RMM) and measurement approaches are available at http://www.cert.org/resilience/rma.html. (See also the report of the Framework track.) Three types of measures or metrics were discussed: implementation, effectiveness, and process performance. Effectiveness measures are the most important and more challenging to implement. The current state of practice is that organizations use implementation metrics.

  Participants noted that the lack of underlying science deprives decision makers of informed alternatives for managing risk. Ideally, metrics should provide answers to such questions as: How many layers of defense-in-depth are really needed? A reference model that enables metrics to be defined in a manner not specific to a single environment is desirable, but when

[6] "Security" here is shorthand for "information security," "computer security," "information assurance," "cyber security," or the union of those problem domains

a metric is applied in a specific context, that context must be taken into consideration.

For cyber resiliency, CERT research is currently focusing on two measures based on DoD Cyber Science & Technology Priority Steering Council guidance: time to restore and effort to restore.

- A MITRE white paper on cyber resiliency metrics, available at https://register.mitre.org/sr/12_2226.pdf. The white paper identifies a wide variety of stakeholders whose decisions could be informed by cyber resiliency metrics, and provides a representative set of metrics, mapped to the Cyber Resiliency Engineering Framework (http://www.mitre.org/work/tech_papers/2012/11_4436/11_4436.pdf).
- The Raytheon paper on cyber defense metrics presented at MILCOM 2010, and cited in the workshop presentation by Suzanne Hassell (https://register.mitre.org/sr/files/hassell.pdf).
- A JHU/APL paper on characterizing cyber resiliency presented at MILCOM 2010, available at http://202.194.20.8/proc/MILCOM2010/papers/p1847-dwivedi.pdf.
- Work on attack-based measurement at MIT/Lincoln Laboratories (http://usacac.army.mil/cac2/cew/repository/presentations/12_Dr%20Scott_MIT_%20Army_Cyber_Symposium_(Publicly_Releasable).pdf) and at Raytheon (http://www.raytheon.com/technology_today/2011_i2/eyeontech_mirror.html).

  In addition, MIT/LL risk analysis technology is incorporated into the FireMon commercial offering (http://www.firemon.com/blog/firemon-acquires-saperix-technologies-whats-next). Other commercial tools also provide some adversary modeling capabilities.

  Measurement can be performed for purposes of assessment or prediction. If the intent is prediction, then an adversary model is needed; the determination of the model should be coordinated with the Framework group.
- A presentation by Roberta Ewart summarizing work presented at the 2006 MORS Symposium (see http://www.promodel.com/pdf/RIST%20Prize.pdf) on a use case-driven analytic approach. See Figure 5. This approach identifies four layers at which goals (and hence target metric values) can be identified, and linked between layers by reference to a use case or scenario. The layers cover most but not all stakeholders. At the campaign, mission, and engagement layers, mission or business process stakeholders think in terms of measures of effectiveness (MOEs), measures of performance (MOPs), and technical performance measures (TPMs). Operations research experts need to think not only at the campaign layer, but also at the mission and engagement layers; similarly, engineers and scientists must think not only at their layers, but also at the mission and engagement layers.

  Thus, the engagement layer is crucial to common understanding. Measurement at that layer can drive questions for science and engineering, particularly for engineering domains (such as resiliency or software) which are less mature than the hardware domain. In general, the community has not made explicit the links between damage to cyber resources and effects on missions; more investment is needed in representing impacts in terms that are meaningful to mission stakeholders.

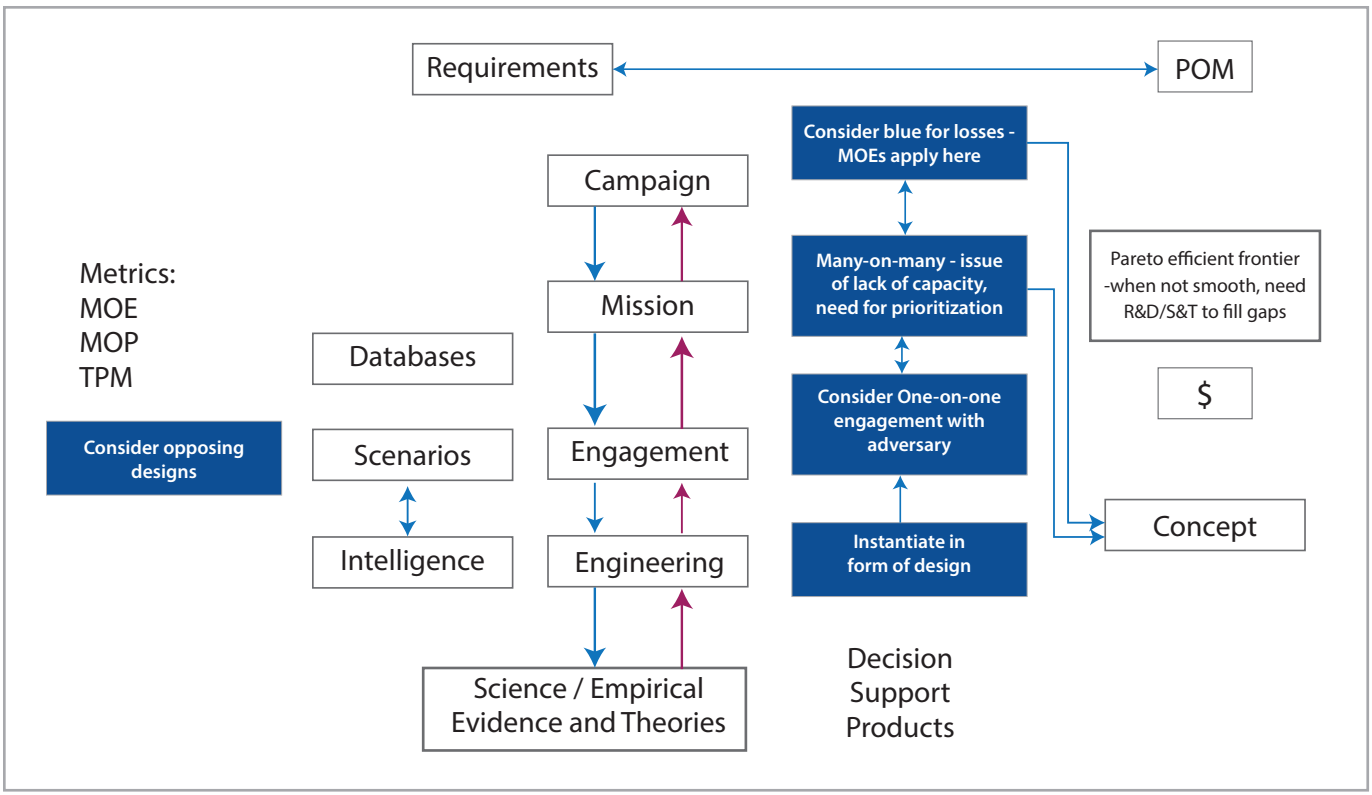

Figure 5. Use Case-Driven Analytic Approach

The resources identified above and in discussion can be organized in terms of a layered representation of stakeholder concerns: The CERT-RMM work applies to the mission; security metrics apply to engagement; work on adversary modeling enables the development of scenarios, which provide the bridge between the engagement and engineering layers; the resiliency work at Raytheon, JHU/APL, and MITRE spans the engagement and engineering layers; and resiliency-related security controls support the development of empirical evidence and theories.

Applying the Use Case-Driven Analytic Approach

Participants identified key questions that must be answered at the Mission, Engagement, and Engineering layers. Those questions can then drive the definition of cyber resiliency metrics.

- Mission / cyber support element of mission
  - o What is critical for the mission to be performed? Cyber resources can be characterized as mission-critical, mission-essential, or supporting. Key challenges include how to represent dependencies that change dynamically, and how to account for non-cyber support.
  - o What if the mission is subverted? What is the mission impact of subversion of cyber support elements? (integrity / accountability)
  - o What if mission cannot be performed? What is the mission impact of loss or degradation of capacity for a given time? (availability / quality of service)
  - o What if the results of cyber support elements of missions cannot be trusted? (integrity)
  - o In the context of the scenarios, what are the mission impacts of the adversary knowing specific information? How does this affect future mission capability? (confidentiality)

  - How can cost / benefits / ROI be represented in the scenarios / model?
- Engagement (one attack at a time)
  - Damage assessment:
    - How far has the adversary gotten?
    - What damage has occurred, and how bad is it?
    - If the damage occurred in the past, how can we discover / account for it?
    - What should we expect in a given set of circumstances? In particular, will the mission be accomplished?
  - Recovery:
    - How well do we understand what constitutes a good state? How do we assess trust? How confident are we in our assessment of trust?
    - How do we know what we need to restore? How quickly can we restore to a known good state?
- Engineering
  - How well have we implemented a specific resilience technique? Or, for a to-be architecture, how well could we? (properties)
  - How well have we defined courses of action, standard operating procedures, defender TTPs, etc.,
    - To take advantage of engineered-in techniques?
    - To compensate for weaknesses?
  - How effective is our implementation? (performance)

For cyber resiliency, adversary modeling can be used to develop a set of scenarios reflecting different adversary motivations, including deny, degrade, destroy, disrupt, and usurp. It is important to distinguish between scenarios in which a mission, network, system, or specific resource is attacked, and scenarios in which the attack occurs in a specific way. Focusing on high-level scenarios supports the definition of global or general metrics, which can then be tailored for specific environments. A key issue for cyber resiliency is the current absence of good models for synergistic or multiple attacks. Another challenge is avoiding "mirroring" – modeling "blue-on-blue" campaigns, missions, and engagements.

Top Principles

Based on the discussion, the following top principles for cyber resiliency metrics and assessment can be identified:

- Do not define, track, or use cyber resiliency metrics as a stovepipe.
  - Build on security metrics experience. Cyber resiliency and security are closely related. Thus, some cyber security metrics can also serve as cyber resiliency metrics, and vice versa In addition, the general process for security metrics can also be applied to the cyber r esiliency problem domain.
  - Integrate resiliency metrics with mission metrics. The relationship between cyber resiliency and mission performance is also close. Thus, the relationship between cyber resiliency metrics and mission Measures of Performance, Measures of Effectiveness, and Technical Performance Measures should be clearly articulated.

  - o Integrate cyber resiliency metrics with organizational metrics and processes. In particular, cyber resiliency can be viewed as a key component of operational resilience, and cyber resiliency practices must be integrated with contingency planning.
- Apply a scenario- or use case-driven analytic approach to define meaningful metrics.
  - o Recognize and account for the role of the threat environment in defining, tracking, and using cyber resiliency metrics. Metrics are meaningful only in the context of assumptions about adversary characteristics and behaviors. Thus, metrics definition must be grounded in adversary modeling.
  - o Identify multiple use cases or scenarios, both to reflect the full range of adversary motivations and to engage – and broker communications among – different stakeholders.
- Metrics need to be justified. Key questions include:
  - o How much does it cost to gather the metric? What does evaluation entail? It was noted that operators currently gather considerable data; if relationships among stakeholders are brokered, existing data could be repurposed to provide metrics that serve multiple stakeholders.
  - o What's the value of the metric? Does the metric provide value? Participants concurred that metrics can have value by improving understanding, as well as by providing prediction as appropriate (with clearly understood uncertainty).
  - o What kind of behavior can the metric be used to drive? Participants concurred that cyber resiliency metrics cannot and should not attempt to be compliance metrics.

Participants also concurred that

- Metrics can be quantitative, semi-quantitative, or qualitative in form.
- Cyber resilience metrics and assessment processes need a broad scope. Resilience is not just a system property; it can apply to an organization, a mission / business process, a network / system, and information.
- The goal-question-metric approach applies – but must be informed by knowledge of whose goal, and whose questions.
- Multiple stakeholders can be identified for cyber resiliency. However, a focus on a few stakeholder communities is necessary in order to avoid confusion. At a minimum, mission / business process owners and executives, architects and systems engineers, resource owners (with a very strong concern for cost and return on investment), and cyber defenders and system/network operators should be considered.

Hard Problems

Participants identified a variety of hard problems, for which further investigation is needed:

- What metrics can help system operators optimize the mixture of resilience courses of action?
- Metrics need to support mission decisions. In what ways can we map cyber resources to missions, so that we can link cyber resilience metrics to mission MOEs/MOPs?
- How do we quantify cyber as a source of risk to mission?
  - – How do we account for scenarios in which we don't see immediate mission effects (e.g., exfiltration)?

  - Which impacts should be considered first, in posing questions to be answered using metrics?
- Immediate vs. long-term vs. potential?
- Focus on availability, or include integrity & confidentiality?
- Scenarios provide a motivation for defining metrics. What should we include in adversary scenarios?
  - Can we take supply chain off the table for initial analysis?
  - How do we account for non-owned assets?
- How do we assess costs?
  - Costs of what?
- Implementing resiliency techniques / solutions (including remediation)
- Not implementing resiliency techniques / solutions
- Sustainment
  - How to represent costs to different stakeholders?
  - How to represent different types of costs? In particular, how to represent mission impact as a form of cost?
  - Can we perform some sort of retrospective assessment of costs?
- How do we account for uncertainty
  - In our models?
  - In our data?
- How do we validate & verify our approach? Our metrics?
  - What is a large enough sample size?
- How do we know that / to what extent the system is trustworthy?
  - How do we define a baseline?
  - How do we measure change? Particularly given the constant change in our systems?
- How can we leverage work in big data?
  - What data is big?
- Sensor data, fused or correlated
- Historical data about system behavior
  - How do we address privacy problems?
  - How do we address prediction challenges?
- How do we define the boundaries / contexts? How do we know whether we've defined them properly? (Note that some things will be invisible if we define the context improperly.)

Directions for Future Work

Participants identified the following ways to move forward in the area of cyber resiliency metrics and assessment:

- Produce a workshop report, which should be shared within the community of practitioners represented at the workshop and socialized more broadly, to engage a wider community.
- Develop an information sharing and discussion venue for the community of practitioners.
- Build connections with the Operations Research community – particularly with metrics groups.

In addition, MITRE plans to revise and update the Metrics paper; this can provide another venue for information capture and sharing.

Engineering Principles for Cyber Resiliency

Track Chair: Kevin Bingham, National Security Agency
Track Co-chair: Harriet Goldman, The MITRE Corporation

Introduction

This paper summarizes the discussions and results from the Engineering Principles Track of the 2nd Secure and Resilient Architectures Workshop held at MITRE, McLean, VA on June 1, 2012. This material will continue to evolve based on continued contributions from workshop participants and from others in the community that join in the group's activities.

Background

The goal of the Engineering Principles group was to identify resilience engineering principles (and a way to reference them) as a basis for:

- Specifying resiliency requirements, operational concepts, and technical designs and architectures.
- Experimenting and integrating into test environments to understand efficacy and define environmental constraints.
- Measuring the cost-effectiveness of techniques both individually, and in combinations.
- Developing a roadmap for maturation and adoption.

The track group acknowledged the immaturity of this area in terms of collective thought and consensus. As a result, the group first discussed the scope and characterization of engineering principles; then held a brainstorming session to collect individual inputs and lessons learned; next organized, categorized and binned the collected ideas; and concluded with beginning to build agreement on the base principles while recording opinions, descriptions and the gauge of consensus on other principles. The session wrapped up with a brief discussion of how the group can continue the collaboration and mature the output.

18

Outcomes and Key Discussion Points

Overview of Track Outcome

- There was significant interest in this track by workshop participants; the group included a large share of workshop attendees and benefited from their diverse backgrounds (academia, industry, government and FFRDCs).
  - Overwhelming community interest in moving this area forward
  - Good interaction and contributions by all
  - Agreement that many such workshop sessions will be needed to reach consensus and increase usefulness of the material
- Scoping the right level and characteristics for capturing principles was challenging. This track discussed proposed approaches but did not reach agreement on a template or hierarchy.
  - While there was general consensus that we should assume best practice for traditional security and build on that to differentiate resilience principles, there was also an acknowledgment that resilience overlaps with certain core Information Assurance (IA) principles and/or that to achieve resilience it may be necessary to extend IA practice beyond traditional applications.
  - The discussion included consideration of principles that deter attacks in the first place; limit the damage and propagation when they are successful; remediate vulnerabilities and return to a trusted state after an attack; and increase adversaries' uncertainty and cost to act as a deterrent against future attacks.
  - The group reached general agreement that resilience principles should:
    - Address both integrity and availability and rely on IA confidentiality capabilities
    - Be high level with supporting mechanisms
    - Be mission-driven and support mission assurance needs
    - Focus on being "resilient enough" and strive for the basics rather than create a complex unachievable vision
    - Be context driven: mission, threat, technology, and environment
    - Apply to people, processes and technology
    - Be first understood individually, but aim at understanding how to create powerful combinations
  - The group agreed it was too early to debate the optimal way to organize principles as ambitiously proposed at track kick-off.
- Significant progress was made during the brainstorming engineering principles session
  - We collected numerous principles; some high level, some more implementation oriented, and others specific to a component or computing environment context.
  - We clustered similar notions together and collected one-of-a-kind or outliers in a separate grouping to be further evaluated.
  - Given the limited time, we selected two specific areas, non-persistence and unpredictability, to delve into more deeply and explore rationale, experiences and lessons learned.
- Overall the track discussions helped the group gauge the:
  - (Im)maturity of this area
  - Diversity of opinions, stakeholders, and individual vantage points

- Challenges that lie ahead to reach consensus
- Importance of context, description and nomenclature
- Desire to continue to work this area
- Need to build on the output of other tracks (framework and metrics) due to interdependencies

Overall Principles

The group agreed we should leverage the initial work done by NIST (SP-800-27, 2004), which includes some resiliency principles as the starting point. The group reached informal consensus on the following broad engineering principles statements for cyber resilient systems:

- Design to reduce exposure to attack
- Design to reduce persistence of access by the adversary
- Design to reduce adversary's ability to act
- Design to limit the consequences of attack
- Design to minimize common cause failure
- Design to tolerate compromise
- Design to degrade gracefully
- Design to crash early and recover quickly

Other principles discussed as noteworthy and accepted by some were:

- Design for integrity and availability
- Design to be threat independent
- Design to be vulnerability independent (we should be concentrating on consequences and how to mitigate them, not on the next vulnerability)
- Design such that users will not seek to circumvent security and resilience features
- Design with distributed and localized decision support
- Integrated horizontally for resilience of the whole system/mission
- Design components with computational plasticity (alternative functional paths to achieve the computing results)
- Design for simplicity and modularity to change easily/frequently

Operational Principles

The group recognized that there needs to be both architectural and operational principles to address the systems (technology, design, and implementation) as well as the people supporting missions and operations (operational and analytical processes and procedures). The following list represents many of the proposed operational principles:

- Actively look for bad guys in the system
- Leverage cyber intelligence to inform operations
- Operate to reduce adversary's ability to act

- Operate to contain vulnerabilities
- Operate to reconstitute and recover quickly to an acceptable level of trust
- Prioritize operational tactics, techniques, and procedures (TTPs) based on mission assurance needs
- Provide situational awareness of IT systems on which missions depend on
- Balance /coordinate local defense with global defense (Tactical actions with global actions)
- Operate to control/ limit the damage/consequences of attack
- Operate with agility and alternative operational contingencies
- Operate to confuse, deceive and impede the adversary (but not the mission operators)
- Monitor integrity and availability, and respond accordingly
- Train operators to understand cyber impact to mission operations to operate in ways that ensure mission execution success

Supporting Principles

As the group tried to bin the ideas from the brainstorming session, it was clear that there were many considerations that must be addressed to achieve resilience beyond architecture and operations. Many suggested principles offered during the brainstorming session were determined to have great merit but were more detailed than the high-level principles captured above and viewed as a "how to" or technique. Other notions were thought to indirectly support high-level resilience principles (e.g., policy and education). In the future work, the group needs to decide whether these ideas should be captured as: a lower level of principle; techniques used to achieve one or more high-level principles either on their own or in combination with other techniques; or as a pre-requisite(s) or support to a principle. The following list captures many of the ideas captured that require further thought.

- More detailed resilience principles
  - Protection/compliance techniques
  - Non-Persistence characteristics
    - o Be random enough to operate inside the adversary's timeline
    - o Be tricky enough that the adversary has to re-earn entry
    - o Be non-persistent in an unpredictable way
    - o Simple reconstitution to a known good state
    - o Randomly discard services
    - o Control availability of assets or features to when they are needed (Just in time services and capabilities)
  - Unpredictability
    - o Doesn't have to be random to be unpredictable
    - o Enable intervening devices to write unassigned fields in network protocols with random data
    - o Create uncertainty
    - o Just in time randomization compilers
    - o Periodic refresh to randomly selected diverse system
    - o Code should be able to run on different hardware/OS combinations
    - o Enable and use different function paths to make your signature less uniform

    - Move operational functionality to a different platform
    - Change operational TTPs
    - Change information flow paths
    - Creation of just in time mission environments
  - Inscrutability/Obfuscation
  - Deception
  - Isolation/Segmentation
  - Discover/Detection
  - Diversity
  - Alternative Processing
  - Redundancy
  - Integrity (hardware, software, data)
  - Alternative Operations
  - Agility
  - High assurance components

- Other Principles
  - Economics of resilience
  - Cyber security and resilience education and training
  - Research and experimentation with resilience (people, process, and technology) in realistic environments. We need
    - To develop an experimental culture
    - Systems that enable observation and experimentation
    - Better models for natural and adversarial events
    - Cyber testbeds and laboratories, modeling and simulation tools
  - Temporal processing principles
  - Measurement/validation of effectiveness
    - Fuzz Testing: Throw random but well formatted inputs into the system and see what it does.
    - Need to define how to test resilience: Red team penetration testing not enough.
    - Trusted Measurement: Need trusted measurement of device health, credential trust, communication links, log, audit, etc. We cannot improve something unless we can measure it.
    - Faults and compromise are ever present in all systems-the types and amounts can change over time.
  - Mission Assurance Engineering
    - Understand mission requirements and dependencies
    - Map the mission and identify crown jewels (Some assets enable ease of mission while others are required for mission completion)
    - Perform threat/vulnerability analysis (Express the threat-vulnerability pairings that degrade, destroy, modify, or re-purpose those assets)
    - Measure risk
    - Define and apply mitigations
    - Test for effectiveness
    - Operational lessons learned

- Architecture
  - o Every specific architecture, approach, and technology must have a threat model and corresponding security claims so engineers can reason about appropriate use and risk
  - o Incorporate resiliency/IA architecture into system architecture. Cannot bolt this on.
  - o Modular design: Based on open standards; supports simplicity and ease of integration and replacement, can enable capabilities (e.g., mobility); enables accountability by provider and specific engineering requirements
- Integration
  - o "Resilience" must mandate principle of large-scale horizontal integration (i.e. meaningful IT product interoperability and interfaces).
  - o Resilience cannot and will not ever be achieved through single vertical solution – no matter how sophisticated or elegant!
  - o No single vendor will own "resilience"
  - o Open standards to drive innovation and interoperability
- Dynamic vs. Static (proactive vs. reactive) resilience principles
- Risk Management
  - o Risk-based requirements and mitigations
  - o Risk-based operational controls Instrumentation
  - o Risk-based decision support
- Domain knowledge and awareness

Future Work

The group agreed that we have just begun to make headway and that further work evolving the material after these initial notes are distributed will be needed. Options on methods to collaborate in the future – such as email, collaboration sites, conference calls, or small workshops – were discussed along with ways to increase visibility and increase community involvement.