MITRE TECHNICAL REPORT MTR120407

MITRE

Cyber Resiliency Assessment: Enabling Architectural Improvement

Project No.: 03127400-GA

Bedford, MA

Deborah Bodeau and Richard Graubart

May 2013

Abstract

Cyber resiliency assessments are intended to identify where, how, and when cyber resiliency techniques can be applied to improve architectural resiliency against advanced cyber threats. This document presents a general process for architectural assessment. The process can be applied to an operational or as-is architecture, to identify first steps or quick wins for improving resilience against advanced cyber threats. The process can also be applied to a notional or to-be architecture, to identify opportunities to provide greater and more cost-effective resilience, and/or to support the development of a cyber resiliency improvement roadmap. The process is supported by assessment scales and questions. Because the set of cyber resiliency techniques continues to evolve, detailed discussion of selected techniques, including POET considerations, is provided.

This page intentionally left blank.

Executive Summary

Cyber resiliency assessments are intended to identify where, how, and when cyber resiliency techniques can be applied to improve architectural resiliency against advanced cyber threats. This document presents a general approach for assessing cyber resiliency and developing recommendations for architectural evolution and process improvement to make more effective use of cyber resiliency practices. The focus is on resiliency assessment for a family of systems, common infrastructure, mission/business segment, or system-of-systems. However, the approach can also be applied to individual systems, services, or components.

The approach can be applied to an operational or as-is architecture, in which case the emphasis may be on "low-hanging fruit" or opportunities for near-term and high-leverage improvements, using a few cyber resiliency techniques. A set of general recommendations provides a starting point for identifying such opportunities. The approach can also be applied to a notional or to-be architecture, in which case the assessment may look at the full set of cyber resiliency techniques, and at ensuring that possible solutions in the mid- and long-term can be integrated into the architecture.

A cyber resiliency assessment requires a structured representation of the problem domain and solution space, so that the scope of the assessment can be clearly defined. The approach uses and extends the Cyber Resiliency Engineering Framework, augmenting the framework with sub-objectives that provide a clearer link between cyber resiliency objectives and techniques, and providing more detailed discussion of the cyber resiliency techniques. Also discussed are applicability of techniques to layers (in a notional layered architecture) and relative maturity of solutions.

Notional examples of cyber resiliency assessments include

- An assessment focused on how Analytic Monitoring capabilities could be improved;
- An analysis of alternatives (AoA) with respect to cyber resiliency objectives and techniques; and
- A comprehensive analysis to support development of a roadmap for improving cyber resiliency.

It is not feasible for organizations to use all of the resiliency techniques. Therefore, POET (political, operational, economic, and technical) considerations for cyber resiliency techniques are identified. It is also not feasible to apply any resiliency technique pervasively, for example due to economic considerations, because implementations of cyber resiliency techniques vary in maturity across different architectural layers, and because some implementations are intended to be used only in strategically chosen locations in a system, common infrastructure, or system-of-systems. Both to serve as a starting point for recommendations and to make more easily understood what the cyber resiliency techniques include, examples are provided of possible solutions that could be integrated in the near-, mid-, and long-term.

Table of Contents

List of Figures

List of Tables

1 Introduction

Missions, business functions, organizations, and nations are increasingly dependent on cyberspace. Attacks in cyberspace are no longer limited to simple (albeit significantly harmful) discrete events such as the spread of a virus or worm, or a denial-of-service attack against an organization. Campaigns are waged by the advanced persistent threat, as illustrated in Figure 1 [1] [2]. Campaigns involve stealthy, persistent, and sophisticated activities, to establish a foothold in organizational systems, maintain that foothold and extend the set of resources the adversary controls, and exfiltrate sensitive information or disrupt operations.

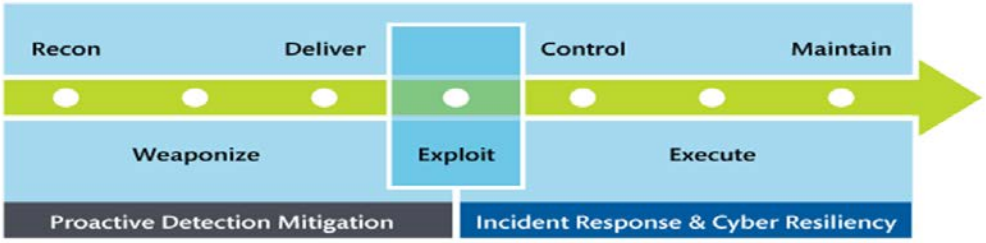

Figure 1. Structure of a Cyber Campaign

Therefore, architecture and systems engineering must be based on the assumption that systems or components have been compromised (or contain undiscovered vulnerabilities that could lead to undetected compromises), and that missions and business functions must continue to operate in the presence of compromise. A growing number of technologies and architectural practices can be used to improve resilience in the face of cyber threats. However, these improvements come with costs as well as benefits. Cyber resiliency assessments are intended to identify where, how, and when cyber resiliency techniques can be applied to improve architectural resiliency in a cost-effective way.

Architectural resiliency is the ability of an architecture[1] – for an enterprise, a mission / business segment, a system-of-systems, a family of systems, or an individual system or component – to enable missions (including cyber defense missions) to anticipate, withstand, recover from, and evolve to address more effectively, cyber-domain attacks. Applying cyber resiliency techniques involves the time-phased integration into architectures of solutions that combine technologies, products, and processes.

Note that cyber resiliency engineering is a relatively young sub-discipline of systems engineering, and approaches to improving resiliency are currently the subject of active research. Thus, the results of resiliency assessments should not be viewed as prescriptive, but instead as recommended ways to start improving cyber resiliency, and to monitor the effectiveness of changes in technologies or processes. There is no "right" or "best" way forward, but there are ways forward, and those ways are found through cyber resiliency assessments.

[1] The term "architecture" can be applied to a range of levels of specificity, and can mean an Enterprise Architecture (EA), the architecture of a mission or business process, the architecture of a mission/business segment (i.e., the set of cyber resources that supports a mission or business process), a system architecture, or the architecture of a product or component. In general, an architecture identifies architectural elements (e.g., systems, services, and common infrastructures for an enterprise architecture or a mission/business segment architecture; components and interfaces for a system architecture), information flows (e.g., transactional flows of mission/business information, control flows of instructions), and functional dependencies.

This document presents a general approach for assessing cyber resiliency and developing recommendations for architectural evolution and process improvement[2] to make more effective use of cyber resiliency practices. The focus is on resiliency assessment for a family of systems, common infrastructure, mission/business segment, or system-of-systems.[3] The approach consists of three steps:

- Determine the scope of, and prepare for, the assessment.
- Assess the architecture.
- Develop specific recommendations.

The approach can be applied to an operational or as-is architecture, in which case the emphasis may be on "low-hanging fruit" or opportunities for near-term and high-leverage improvements, using a few cyber resiliency techniques. A set of general recommendations provides a starting point for identifying such opportunities. The approach can also be applied to a notional or to-be architecture, in which case the assessment may look at the full set of cyber resiliency techniques, and at ensuring that possible solutions in the mid- and long-term can be integrated into the architecture.

An assessment can be very high-level (assessment of cyber resiliency objectives), more detailed (assessment of cyber resiliency sub-objectives and of how well relevant cyber resiliency techniques are applied), or very detailed (assessment of resiliency-enhancing activities, which support sub-objectives and are identified with cyber resiliency techniques; determination of how different solutions would change the assessments of activities, techniques, and sub-objectives). This document focuses on high-level and more detailed assessments, providing assessment value scales for objectives and sub-objectives, as well as general value scales applicable to all cyber resiliency techniques with examples of technique-specific values scales for a few techniques. To illustrate the level of a very detailed assessment, a few examples of value scales for resiliency-related activities are also provided.

1.1 Relationship to Other Assessment Approaches

Architectural assessment approaches can focus on:

- Properties: To what extent does the architecture achieve cyber resiliency objectives, or how effectively does it incorporate cyber resiliency techniques?
- Processes: How well does an organization execute processes related to cyber resiliency?
- Performance: How quickly, how correctly, and with what degree of confidence can cyber-supported mission or business functions, and supporting cyber defender activities, be performed in the presence of attacks by advanced adversaries?

The approach presented in this document focuses on assessment of architectural properties, for which improvement involves primarily adoption of new or more effective use of existing technologies. Such an assessment can be performed for an as-is (operational) or to-be (notional)

[2] Particularly in the near term, improvements can be made to administrator Standard Operating Procedures (SOPs) and cyber defender tactics, techniques, and procedures (TTPs) that use existing technology more effectively to improve cyber resiliency.
[3] The overall approach also applies to systems and components; however, the supporting tables in Appendix C must be tailored.

architecture. For an as-is architecture, the assessment[4] can be informed by performance metrics [3] or complemented by assessments of processes and performance [4] [5]. The Realignment cyber resiliency technique, the Transform and Re-Architect objectives, and the Evolve goal involve primarily process improvement. Thus, the approach presented in this document provides only limited coverage of these aspects of cyber resiliency.

1.2 Overview of This Document

A cyber resiliency assessment requires a structured representation of the problem domain and solution space, so that the scope of the assessment can be clearly defined. Section 2 provides background on the Cyber Resiliency Engineering Framework, on which the architectural resiliency assessment process is based [6]. The initial framework is augmented, with sub-objectives that provide a clearer link between cyber resiliency objectives and techniques. Section 3 provides notional examples of cyber resiliency assessments.

Sections 4-6 present the general assessment process. The approach described in this document can be applied to all cyber resiliency objectives and techniques, using the general assessment scales presented in Section 4 as a starting point. Section 5 provides recommendations for areas in which near-term improvements can be sought, as well as general recommendations for applying cyber resiliency techniques. These recommendations can serve as a starting point for developing specific recommendations for a given architecture.

Appendix A provides a more detailed mapping between cyber resiliency objectives and techniques. Appendix B provides an initial set of questions that can be used when assessing an as-is architecture. Appendix C provides scales that can be used in a cyber resiliency assessment; values are defined for cyber resiliency objectives, sub-objectives, and representative examples of cyber resiliency techniques and cyber defender activities. Appendix D amplifies the description of cyber resiliency techniques from [6], to support development of recommendations. It is not feasible to apply any resiliency technique pervasively, for example due to economic considerations, because implementations of cyber resiliency techniques vary in maturity across different architectural layers, and because some implementations are intended to be used only in strategically chosen locations in a system, common infrastructure, or system-of-systems. Appendix E identifies POET (political, operational, economic, and technical) considerations for cyber resiliency techniques, and Appendix F provides examples of possible solutions that could be integrated in the near-, mid-, and long-term.

[4] This document defines a set of qualitative metrics to be used in an architectural assessment. Quantitative metrics, such as those included in [3], can be used as supporting evidence for qualitative assessments. The Federal Enterprise Architecture Assessment Framework [165] provides an example of how quantitative metrics can be used to support qualitative assessments: For some key performance indicators (e.g., Scope of Completion), achieving a level of performance entails demonstrating that values of specified metrics are at or above specified levels.

2 Background

This section provides background on the Cyber Resiliency Engineering Framework, which can be used to structure analysis during an assessment. As shown in Figure 2, the framework organizes the cyber resiliency domain into a set of goals, objectives, and techniques. Goals are high-level statements of intended outcomes. Objectives are more specific statements of intended outcomes, expressed so as to facilitate assessment; an objective can be identified with a single goal but may support achieving multiple goals.

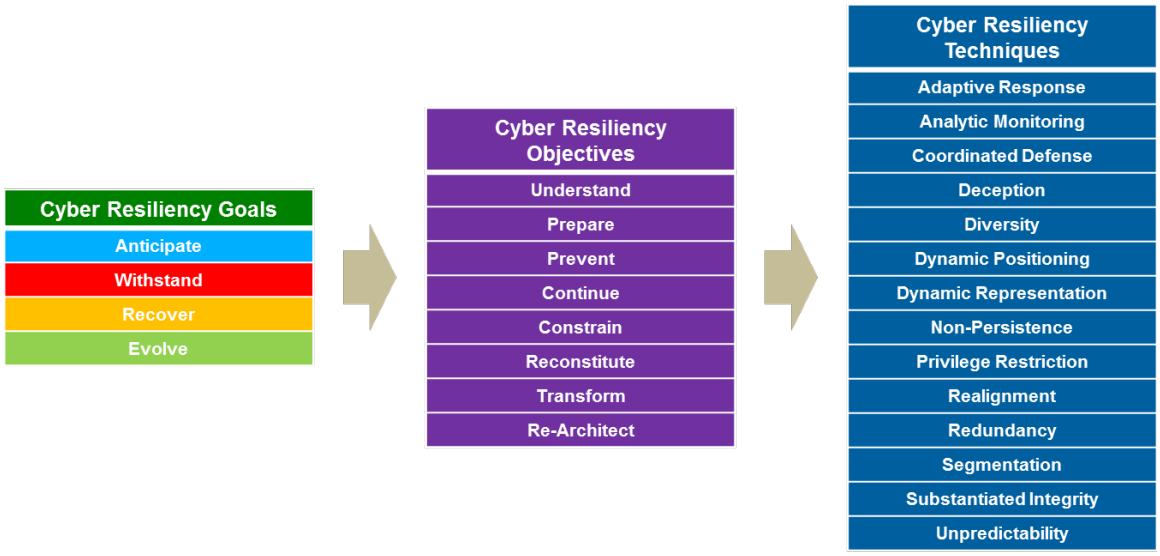

Figure 2. Cyber Resiliency Goals, Objectives, and Techniques

In the context of cyber resiliency, techniques are approaches to achieving one or more cyber resiliency objectives that are applied to the architecture or design of mission/business functions and the cyber resources that support them. Cyber resiliency techniques are selectively applied to an architecture or to the design of mission/business functions and the cyber resources that support them to achieve objectives. A given technique usually supports multiple objectives, but may be unique to a single objective.

2.1 Mapping Techniques to Objectives Using Sub-Objectives

Cyber resiliency objectives and techniques are defined at a high level, which makes the relationships among them difficult to discern. Therefore, the next level of detail in the framework consists of sub-objectives and capabilities. The relationship between a cyber resiliency objective, and a technique that could be applied to meet that objective more fully and effectively, is based on the capabilities or cyber defender actions it enables. Mappings are provided in Appendix A. The sub-objectives and techniques are shown in Table 1.

Table 1. Cyber Resiliency Objectives and Sub-Objectives Enabled by Techniques

Objective	Sub-Objective	Techniques
Understand: *maintain useful representations of mission/business cyber dependencies, and of the status of cyber resources with respect to possible adversary activities*	Understand adversaries	Analytic Monitoring
		Deception
	Understand mission or business function dependencies on cyber resources *and* Understand the functional dependencies among cyber resources	Dynamic Representation
		Realignment
		Coordinated Defense
		Privilege Restriction
	Understand the status of resources with respect to adversary activities	Adaptive Response
		Analytic Monitoring
		Dynamic Positioning
		Dynamic Representation
		Substantiated Integrity
Prepare: *maintain a set of realistic cyber courses of action that address predicted or anticipated cyber attacks*	Create and maintain cyber courses of action	Coordinated Defense
	Maintain the resources needed to accomplish cyber courses of action	Coordinated Defense
	Validate the realism of cyber courses of action	Coordinated Defense
		Dynamic Representation
Prevent: *preclude successful execution of an attack on a set of cyber resources*	Harden resources based on adversary capabilities	Coordinated Defense
	Deflect adversary actions	Deception
	Dissuade / deter adversaries by increasing the adversary's costs	Diversity
		Privilege Restriction
		Segmentation
		Unpredictability
	Dissuade / deter adversaries by increasing the adversary's risks	Analytic Monitoring
		Deception
	Deter attacks by limiting the adversary's perceived benefits	Deception
		Non-Persistence
Continue: *maximize the duration and viability of essential mission/business functions during an attack*	Maintain functioning	Adaptive Response
		Diversity
		Coordinated Defense
	Ensure that functioning is correct	Substantiated Integrity
	Extend the surface an adversary must attack to be successful	Privilege Restriction
		Non-Persistence
		Unpredictability
Constrain: *limit damage from an adversary's attacks*	Isolate resources to preclude or limit adversary access	Segmentation
	Move resources to preclude adversary access	Dynamic Positioning
		Realignment
	Change or remove resources to limit or preclude adversary access	Non-Persistence
		Privilege Restriction
		Adaptive Response

Objective	Sub-Objective	Techniques
Reconstitute: *redeploy cyber resources to provide as complete a set of mission / business functionality as possible subsequent to a successful attack*	Maintain deployable / redeployable resources	Redundancy
	Restore functionality	Adaptive Response
		Coordinated Defense
	Validate functionality	Substantiated Integrity
Transform: *change aspects of organizational behavior in response to prior, current, or prospective adversary attacks*	Identify unnecessary dependencies	Realignment
	Adapt systems and mission / business processes to mitigate risks	Realignment
Re-Architect: *modify architectures for improved resiliency*	Address predicted long-term changes in adversary capabilities, intent, and/or targeting	Supporting (Systems and Systems-of-Systems [SoS] Engineering, drawing from Analytic Monitoring)
	Apply cyber resiliency practices cost-effectively	Supporting (Systems and SoS Engineering)
	Incorporate emerging technologies in ways that improve (or at least do not degrade) cyber resiliency	Supporting (Systems and SoS Engineering)

2.2 Application Domains for Cyber Resiliency Techniques

Cyber resiliency techniques can be applied at different domains (layers in a notional layered architecture), as indicated in Table 2. Effective application of cyber resiliency techniques to different layers leverages approaches from the broader disciplines of fault-tolerant computing, network resilience, and system resilience using redundancy for backup, failover, and recovery.

Table 2. Application Domains for Cyber Resiliency Techniques

Application Domain / Layer	Examples	Related Resilience Approaches
Hardware/firmware	FPGA, MPSoC, general and specialized processors, embedded firmware	Fault-tolerant hardware
Networking/communications	Communications media, networking protocols	Network resilience, especially using redundancy
System/network component	Firewalls, servers, thin-clients	Fault-tolerant design
Mobile system/network component	Laptops, tablets, smartphones, PDAs (transiently or intermittently part of a system or network)	Fault-tolerant design (especially tolerance of drops in connectivity)
Operating system	General-purpose OS, RTOS	Fault-tolerant design
Cloud, virtualization, and/or middleware infrastructure	VMM, hypervisor, SOA infrastructure / shared services	Fault-tolerant design; middleware for predictable and load-balanced service
Mission / business function application / service	Tailored DBMS, workflow management software; specialized mission applications	Fault-tolerant design
Software	Software running on system/network components (including OS, cloud, virtualization, middleware, DBMSs, applications, services)	Fault-tolerant design

Application Domain / Layer	Examples	Related Resilience Approaches
Information stores	Databases, knowledge bases, unstructured collections ("big data"[5])	System resilience using redundancy for backup, failover, and restore/rollback
Information streams / feeds	RSS feeds, Twitter, instant messaging / chat, video feeds	Network resilience, especially using redundancy
Systems[6]	Integrated sets of the foregoing, within a single administrative or management span of control.	System resilience using redundancy for backup, failover, and restore
Systems-of-systems[7]	Sets of systems under multiple spans of control, which interoperate to support a given mission or set of missions. Within an organization, a system-of-systems is at Tier 2 in the Risk Management Hierarchy; however, a system-of-systems can span multiple organizations.	System resilience using redundancy for backup, failover, and restore; network resilience using redundancy for alternate communications paths

2.3 Overarching Questions for a Cyber Resiliency Assessment

As illustrated in Figure 3, the questions that a cyber resiliency assessment could answer can be mapped to a representation of the problem space (goals, objectives, and sub-objectives) and the solution space (techniques, activities, and solutions).

The cyber resiliency goals of Anticipate, Withstand, Recover, and Evolve are defined to be consistent with frameworks from resilience engineering, national preparedness, and resilient networks, among others. However, goals are defined at too high a level for the extent to which they are achieved to be assessed directly; this is why the framework defines objectives. The assessment of how well goals are met is derived from the assessment of how well objectives (or sub-objectives) are achieved.

[5] "Big data refers to large datasets that are challenging to store, search, share, visualize, and analyze." Examples include "User and machine-generated content through social media, web and software logs, cameras, information-sensing mobile devices, aerial sensory technologies, and genomics." [165]

[6] A system is at Tier 3 in the Risk Management Hierarchy [57].

[7] The Defense Acquisition Guidebook [164] defines a system-of-systems (SoS) as "a set or arrangement of systems that results from independent systems integrated into a larger system that delivers unique capabilities." The DoD Systems Engineering Guide for Systems of Systems [165] identifies four types of systems-of-systems: Virtual, Collaborative, Acknowledged, and Directed.

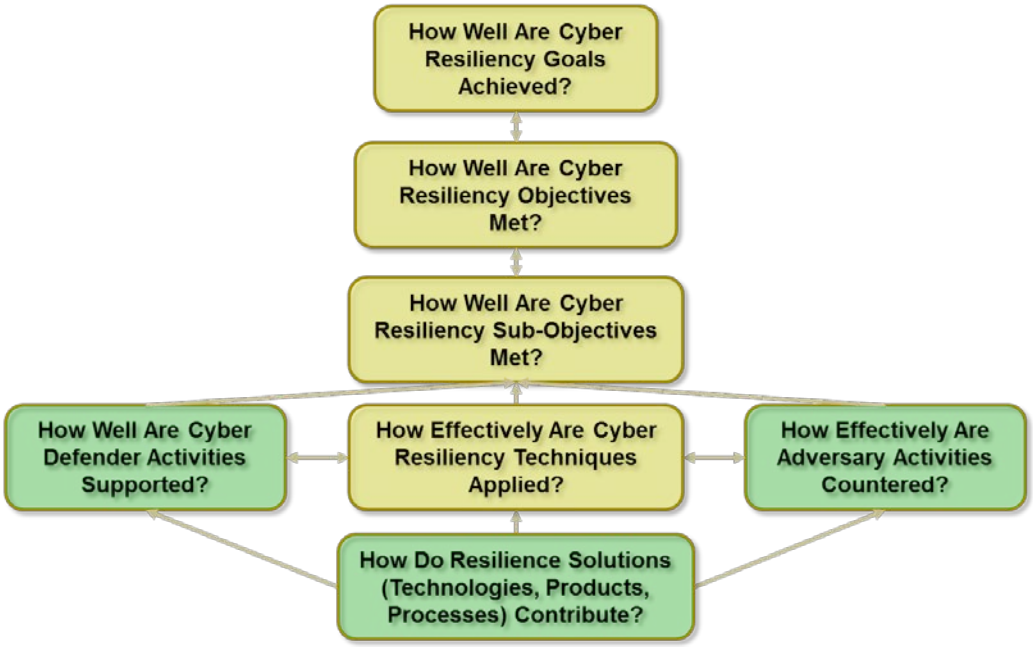

Figure 3. Key Questions to Be Considered in an Assessment

Objectives and techniques can be assessed directly, to provide an overall sense of architectural resiliency. However, to develop recommendations for techniques in which to invest, an assessment should include at least sub-objectives, and may also include relevant cyber defender activities.[8] (Assessment of how effectively cyber adversary activities are countered is a topic for future research.) If specific solutions are part of the scope of the assessment, their contributions should also be assessed.

An assessment can be very high-level (assessment of cyber resiliency objectives), more detailed (assessment of cyber resiliency sub-objectives and of how well relevant cyber resiliency techniques are applied), or very detailed (assessment of resiliency-enhancing activities, which support sub-objectives and are identified with cyber resiliency techniques; determination of how different solutions would change the assessments of activities, techniques, and sub-objectives). This document focuses on high-level and more detailed assessments, providing assessment value scales for objectives and sub-objectives, as well as general value scales applicable to all cyber resiliency techniques with examples of technique-specific values scales for a few techniques. To illustrate the level of a very detailed assessment, a few examples of value scales for resiliency-related activities are also provided.

[8] Examples of sub-objectives include Understand Adversaries, Create and Maintain Cyber Courses of Action (CCoAs), and Validate Functionality. Examples of cyber defender activities include Dynamically Relocate Sensors, Track Effectiveness of CCoA and Adapt as Necessary, and Validate Data Provenance. See Appendix A for a full listing of sub-objectives and activities.

3 Notional Examples of Cyber Resiliency Assessments

This section provides notional[9] examples of ways the cyber resiliency assessment process and material in this document could be applied. An assessment can range from light-weight to comprehensive. The level of effort depends on the purpose and scope of the assessment.

For example, an assessment could focus on how Analytic Monitoring capabilities could be improved. Figure 4 illustrates how the material in this document can be used to produce recommendations:

- The differentiating factors for Analytic Monitoring (identified in Table 7, Key Differentiators Between Levels for Cyber Resiliency Techniques) were used, together with the general definitions of levels (Table 8, Representative Reasons for Restricting Consideration of Cyber Resiliency Technique) to create descriptions of the different levels (Table 16, Definitions of Assessment Levels and General Recommendations for Analytic Monitoring), accompanied by general recommendations.[10]
- Existing capabilities are assessed, and the general recommendations from Tables 11 (General Recommendations for Applying Cyber Resiliency Techniques) and 16 are tailored, as illustrated. Tailoring is based on POET considerations (Table 39, POET Considerations for Analytic Monitoring), as well as solutions that could be phased in over time (Analytic Monitoring line of Table 52, Representative Examples of Cyber Resiliency Mechanisms).

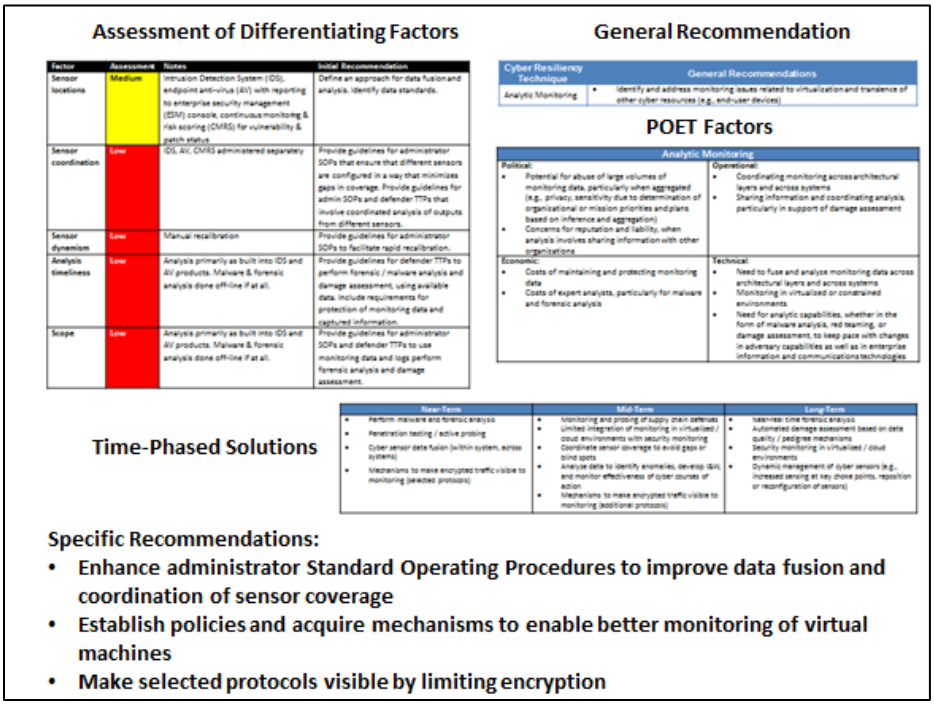

Figure 4. Notional Assessment and Recommendations for Analytic Monitoring

[9] No specific program, network, mission/business segment, or data center should be inferred from these examples.

[10] Appendix C provides value scales and corresponding general recommendations for a representative set of cyber resiliency techniques. When a cyber resiliency assessment includes a technique not represented in Appendix C, Tables 7 and 8 can be used to construct additional tables, following the examples in Appendix C.

A light-weight process could help the Program Manager (PM) or owner of a family of systems identify architectural alternatives for improving cyber resiliency, or include a cyber resiliency component in an analysis of alternatives (AoA). Such a process could consist of an interview with the PM/owner and review of high-level documentation (e.g., concept of operations, requirements, architecture, design, and continuity of operations planning documents), analysis taking one to two weeks, and presentation of results in briefing form. A light-weight process could go down to the level of cyber resiliency objectives or sub-objectives, or could focus on a few cyber resiliency techniques. Appendix D can be used to gain more insight into cyber resiliency techniques.

Figure 5 illustrates an AoA that looks at four alternatives with respect to

- The cyber resiliency objectives, using Table 13 (Assessment Scales for Cyber Resiliency Objectives).
- The level of architectural flexibility (i.e., ability to use, support, or otherwise accommodate) with respect to the different cyber resiliency techniques, using Table 8 (Levels of Implementation).

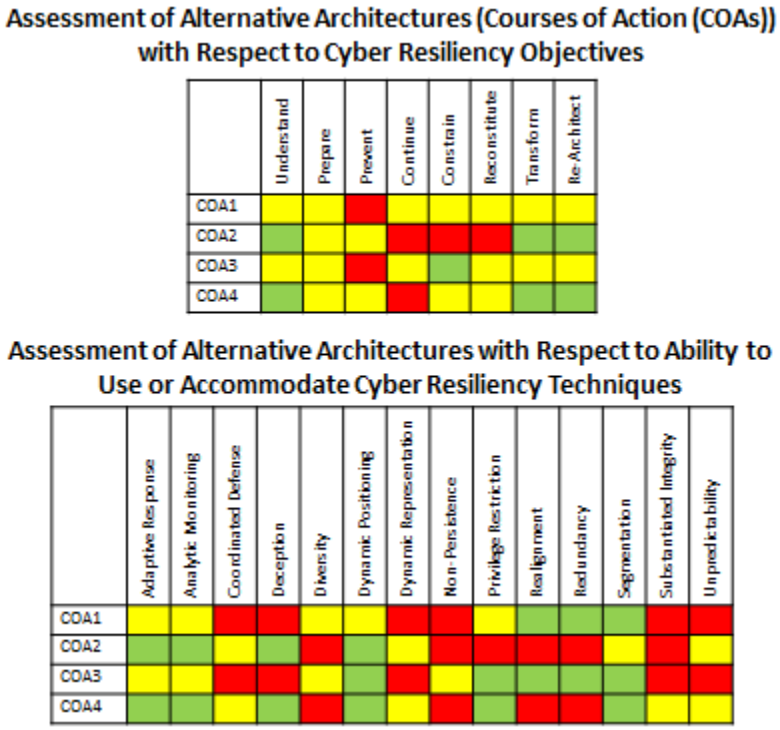

Figure 5. Notional Assessment of Alternatives

A fairly comprehensive process could help cyber defenders and PMs for a system-of-systems supporting multiple missions develop a roadmap for improving cyber resiliency. Such a process could involve analysis taking a month or more, and

- Include interviews with multiple stakeholders (e.g., mission owners, PMs, commanders of cyber defense operations) and subject matter experts (e.g., systems engineers, system administrators, cyber defenders, systems architects and engineers). The questionnaire in Appendix B can serve as a starting point for such interviews. See Table 4 (Possible Stakeholders and Subject Matter Experts to Interview).

- Include review of a large body of documentation. See Table 5 (Possible Source Documents).
- Result in a detailed report.

Material in this document could be used in the following ways:

- Assess cyber resiliency sub-objectives using Table 14 (Assessment Scales for Cyber Resiliency Sub-Objectives). Summarize results by rolling up assessments of sub-objectives into assessments of objectives, consistent with Table 13 and taking into consideration the relative importance stakeholders place on sub-objectives.
- For sub-objectives rated Low, use Table 12 (Detailed Mapping of Cyber Resiliency Techniques to Sub-Objectives and Objectives) identify cyber resiliency techniques to assess in more detail, and the capabilities those techniques should enable. (Time and effort permitting, a second pass could be performed, for cyber resiliency techniques and capabilities related to sub-objectives rated Medium.)
- Further refine the set of cyber resiliency techniques to consider, using Table 3 (Representative Reasons for Restricting Consideration of Cyber Resiliency Techniques).
- Assess the level of architectural flexibility (i.e., ability to use, support, or otherwise accommodate) with respect to the different cyber resiliency techniques. Use Tables 15-20 for Adaptive Response, Analytic Monitoring, Coordinated Defense, Diversity, Privilege Restriction, and Redundancy. Use Table 7 (Key Differentiators Between Levels for Cyber Resiliency Techniques) and Table 8 (Levels of Implementation) to provide more specific characterizations of the levels for other techniques.
- Develop recommendations.
    - Identify POET considerations, using the relevant tables in Appendix E as a starting point.
    - Develop an initial set of recommendations for architectural evolution and process improvement, using the recommendations in Tables 15-20 as a model.
    - Identify specific mechanisms that could be phased in, using Table 52. Tailor, based on POET considerations[11]; augment using the discussion and tables (Application and Maturity) in Appendix D.

[11] Note that metrics related to cost can be developed following the guidance in [6].

4 Determine the Scope and Plan for the Assessment

Planning an assessment involves

- Determining the purpose of the assessment;
- Determining the scope of the assessment; and
- Identifying key stakeholders and sources of information.

4.1 Determine the Purpose

The purpose of an assessment is defined by the questions it is intended to answer and/or the decisions it is intended to support. Those questions or decisions should initially be expressed in stakeholder terms, rather than in terms of cyber resiliency; they can then be translated into the terminology of the cyber resiliency framework, as illustrated in Figure 2 in Section 2.3. The purpose of an assessment is defined based on discussion with the individual decision-maker or set of decision-makers whom it is intended to support. Examples include:

- Establishing a baseline representation of the current posture of an as-implemented architecture, and developing recommendations for first steps to improve cyber resiliency. The baseline representation can consist of an assessment with respect to meeting relevant cyber resiliency goals, achieving relevant cyber resilience objectives, and/or applying relevant cyber resiliency techniques. (See the discussion of scope, below, for how relevance is determined.)
- Assessing a notional, to-be, or as-implemented architecture's ability to achieve relevant cyber resiliency objectives or apply cyber resiliency techniques, and developing a roadmap for effective integration of cyber resiliency technologies and processes. The assessment can identify areas in which techniques can easily be applied, as well as gaps in or impediments to cyber resiliency techniques.
- Performing an analysis of alternatives (AoA) for different notional or to-be architectures, with respect to each architecture's ability to achieve relevant cyber resiliency objectives or apply cyber resiliency techniques.

4.2 Determine the Scope

The scope of a cyber resiliency architectural assessment consists of:

- The architecture[12] to be assessed. A cyber resiliency assessment typically is performed for a family of systems, a common infrastructure or shared service, a mission/business segment, or a system-of-systems, but can also be performed for an individual system or even a hardware or software component. For assessments of architectures above the component level, the technical architectural description is complemented by a description of the missions and/or business processes supported by the architecture.[13]

[12] The term "architecture" refers to a description of components, their internal structures, their relationships, and the principles and guidelines governing their design and evolution over time [165]. An architecture can be represented in multiple ways [164], and can cite technical standards as part of the principles and guidelines.

[13] That is, a technical architecture can be viewed as part of a mission architecture. See, for example, [164]. Alternately, mission threads can be described [165].

- The cyber resiliency objectives and techniques to be considered. An assessment can include all objectives and techniques. However, an assessment can be restricted to a subset of objectives, based on stakeholder priorities, and/or to specific techniques. An assessment can even be restricted to a set of cyber defender activities. These restrictions are determined by the purpose of the assessment.

The scope of an assessment is also determined in part by the architectural layers that are included in the architecture. For example, an assessment of a network (an example of a common infrastructure) can disregard information diversity.

Meeting cyber resiliency objectives depends on processes, procedures, and governance. Therefore, an assessment of a notional or to-be architecture will either omit assessment of how well objectives are met, or will ask how well objectives can be expected to be achieved, based on assumptions (which must be stated) about the concept of operations.

For a variety of reasons, it is not feasible for organizations to use all of the resiliency techniques. The determination of which to apply is part of an overall risk management process. POET considerations, as discussed in Appendix E, help make that determination. Table 3 provides representative examples of reasons for excluding techniques from consideration or restricting specific techniques to be considered in developing recommendations.

Table 3. Representative Reasons for Restricting Consideration of Cyber Resiliency Techniques

Technique	Representative Reasons for Restricting Consideration
Adaptive Response	Liability concerns (e.g., responses that violate SLAs, cause collateral damage)
Analytic Monitoring	Policy concerns related to collecting, aggregating, and retaining data (e.g., sensitivity / classification, privacy)
Coordinated Defense	Governance and CONOPS issues (e.g., overlapping or incompletely defined roles and responsibilities, no clear responsibility for defining cyber courses of action)
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 which 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)

Alternately, an assessment can focus on a few techniques, based on organizational priorities. See Section 6 for recommended priorities to achieve quick wins in cyber resiliency improvement.

4.3 Identify Information Sources

A cyber resiliency assessment can draw upon multiple sources of information, including stakeholders and documentation.

4.3.1 Stakeholders

The determination of which resiliency techniques are needed is largely driven by the needs of the stakeholders. Different stakeholders have different priorities. To best understand the priorities of the stakeholders requires interviewing them. Table 4 identifies possible stakeholders and subject matter experts (SMEs) who might be interviewed. The set of stakeholders and SMEs who are interviewed for an assessment depend on the assessment's scope. The expertise of three types of SMEs is particularly important:

- Architects and systems engineers. These SMEs can provide an understanding of the architecture, as well as the POET considerations that constrain the solution space.
- Cyber defenders (i.e., staff whose role makes them responsible for defending cyber resources against attack, and thus for responding to detected or suspected incidents). Cyber defenders can be characterized[14] as
    - Tactical or line-level management, e.g., staff in a Tier 3 Network Operations and Security Center (NOSC), a Cyber Security Operations Center (CSOC), or a Computer Emergency Response Team (CERT); system administrators or managers for whom response to cyber incidents is part of defined responsibilities.
    - Operational or mid-level management, e.g., staff in a Tier 2 Cyber Operations Center (CyOC).
    - Strategic or high-level (enterprise-level) management, e.g., staff in a Joint Operations Center (JOC).

  For purposes of architectural assessment, tactical and operational cyber defenders are important SMEs, since they can discuss cyber courses of action (CCoAs) and how – and how well – the architecture enables them to fulfill their responsibilities by executing CCoAs. However, it must be noted that some organizations do not have cyber defenders, or outsource cyber defense functions.
- Program Managers and information technology (IT) or information and communications technology (ICT) providers and operators (e.g., the manager of a fixed-site facility that provides computing resources to multiple missions or users, the provider of a common infrastructure or set of shared services; system administrators or managers). Providers and operators can identify performance criteria and requirements, particularly those captured in service level agreements (SLAs). In addition, administrative staff can provide

[14] The DoD has defined three tiers of Computer Network Defense (CND), which roughly correspond to the three tiers in the multi-tiered approach to risk management in NIST SP 800-39. For some organizations, the tiers are collapsed, so that (for example) operational and tactical decisions are made by the same set of cyber defenders.

insight into whether and how standard operating procedures (SOPs) apply cyber resiliency techniques.

Table 4. Possible Stakeholders and Subject Matter Experts to Interview

Role	Information to Obtain
Mission Owner	Mission priorities – what tasks are mission-essential, mission-critical, or supportive; relative priority of near-term vs. long-term mission capabilities. Used to determine relative importance of cyber resources, relative importance of cyber resiliency goals and objectives. Usually not interviewed; priorities are reflected in system / program requirements and in Mission / Business Impact Analyses.
Cyber Defender	How – and how well – the architecture enables cyber defenders to fulfill their responsibilities. (See Appendix A.)
Program Manager	Relative priorities of cyber resources and cyber resiliency goals and objectives, based on the missions the program supports, the relative priorities of near-term vs. long-term capabilities for those missions, and the criticality of cyber resources to those missions.
IT/ICT Provider (e.g., Datacenter Manager)	Relative importance of different capabilities or services. Capabilities related to continuity of operations and execution of contingency plans.
Architect / Systems Engineer	Current and future architecture. POET considerations, particularly technical constraints.

4.3.2 Documentation

Table 5 identifies possible source documents for a cyber resiliency assessment. The source documents consulted depend on the scope of the assessment, and on the life-cycle stage (for a component, system, family of systems, common infrastructure, or set of shared services).

Table 5. Possible Source Documents

Source Document	Relevance
Mission Impact Analysis or Business Impact Analysis	Identifies mission (or business process) concerns and priorities. Identifies mission-essential and mission-critical resources. Provides basis for contingency plans.
Contingency Plans, including Business Continuity Plan (BCP), Continuity of Operations Plan (COOP), and Information System Contingency Plan (ISCP) (see NIST SP 800-34 Rev. 1 for descriptions)	Describes how cyber resources and operational processes (e.g., cyber defense, system administration) are used to ensure mission / business continuity under stress. Thus, identifies functions and resources that support cyber resiliency.
Architecture documentation	Describes the architectures of the mission / business segment, system-of-systems, common infrastructure, set of shared services, system, and/or components.[15]
Standard Operating Procedures (SOPs) for system or network administration, and for handling computer incidents	Describes how cyber resources are used to enforce policies and meet SLAs. Describes operational processes for responding to incidents, including which functions / cyber resources are used.
CND plans and procedures, cyber courses of action (CCoAs), or cyber playbooks[16]	Describes processes, procedures, and cyber resources used in those processes for cyber defense.

[15] Note that the architecture must be described in enough detail that it can be analyzed to identify architectural layers, identify defensive techniques, and perform a mapping from techniques to layers to enable assessment of breadth of defense and depth of defense.

[16] A cyber playbook is an encoding of knowledge about how to handle cyber situations ranging from notification of attacks against other organizations to evidence of an adversary campaign. [165]

5 Perform the Assessment

Performing the assessment involves assigning qualitative values ranging from very low to very high, with supporting rationale. As illustrated in Section 3, the specific tables defining the qualitative values will be selected based on the assessment's purpose and scope. Tables defining values for cyber resiliency objectives and sub-objectives are presented in Appendix C.

This section provides general value scales that can be used for all cyber resiliency techniques. However, for some assessments, a more detailed set of definitions may be needed. Appendix C includes examples of more detailed value scales for six techniques: Adaptive Response, Analytic Monitoring, Coordinated Defense, Diversity, Privilege Restriction, and Redundancy. Tables defining values for cyber resiliency techniques can be constructed, using the general structure presented in Section 5.1, and the examples for those six techniques.

For an operational (or "as-is") architecture, the level of implementation can also be assessed, taking into consideration real-world CONOPS, SOPs, and as-deployed mechanisms. In addition, the ability to perform cyber defender activities can be assessed, if that level of detail is consistent with the purpose and scope of the assessment. Section 5.2 provides a general value scale for levels of implementation, and for resilience-related cyber defender activities. Appendix C provides a few examples.

The process of selecting and tailoring value scales, and evaluating with respect to those scales, can be multi-step. For example, in-scope objectives could be evaluated to identify those for which need is greatest; for those objectives, sub-objectives could be evaluated to get a more nuanced identification of improvement needs. The mapping of sub-objectives to resiliency techniques in Appendix A could be used to identify techniques to focus on when developing recommendations. Those techniques that are in scope (based on POET considerations) could then be evaluated.

Documenting the rationale for the assignment of a value is a crucial part of the assessment. The rationale can identify inherent architectural vulnerabilities to attacks by advanced adversaries; gaps in technologies, SOPs, cyber courses of action (CCoAs); and/or policy or governance issues. This provides the foundation for developing recommendations.

5.1 Architectural Flexibility or Capability

For any architecture[17], an assessment can determine how flexible or capable the architecture is with respect to the incorporation and effective application of a resiliency technique. Table 6 provides a means of assessing the relative flexibility on a scale from Very Low to Very High.

As noted earlier, not all resiliency techniques will be employed in a given architecture. Similarly, not all aspects of a technique may be applicable for all architectures. Different factors comprise each of the resiliency techniques. These differentiating factors are identified as in Table 7. Assessment levels can be defined for each factor; assessment of differentiating factors can be useful when supporting an Analysis of Alternatives or developing specific recommendations. For examples, see the tables in Appendix C.2.[18]

[17] Underlying this statement is the assumption that the architecture is sufficiently mature and well that analysis can be performed against it. For architectures that are simply at the conceptual/planning stage it is often the case that there is not sufficient information available to perform a meaningful assessment.

[18] Definitions of assessment levels are given for Low, Medium, and High; Very High is typically a stretch goal. For Very Low, the supporting explanation in the assessment should identify the architecture's limitations.

Table 6. Definitions of Levels for Flexibility

Level	Description
Very High	The architecture **explicitly integrates** a **strategically-chosen** set of components, technologies, and processes to implement the resiliency technique, includes **mechanisms to assess the effectiveness** of those technologies and processes, and **explicitly provides the flexibility** to integrate additional technologies or components as they become available.
High	The architecture **explicitly includes** components, technologies, and processes to implement the resiliency technique, and provides **some flexibility** to integrate additional technologies or components as they become available.
Medium	The architecture **accommodates or includes** components, technologies, and processes to implement the resiliency technique, and provides **some flexibility** to include additional technologies or components as they become available.
Low	The architecture **does not preclude** components, technologies, and processes to implement the resiliency technique, and provides **limited flexibility** to include additional technologies or components as they become available.
Very Low	The architecture **precludes** many components, technologies, and processes that could implement the resiliency technique, and provides **severely limited opportunities** for future consideration of the additional technologies or components as they become available.

Table 7. Key Differentiators Between Levels for Cyber Resiliency Techniques

Cyber Resiliency Technique	Key Differentiators Between Levels
Adaptive Response: take actions in response to indications that an attack is underway based on attack characteristics	• **Breadth of response**: How many different responsive actions does the architecture support? • **Depth of response**: At how many architectural layers, or for how many architectural components, can responsive actions be taken? • **Dynamism**: How quickly can response actions be taken? • **Integration**: How well are other resiliency technologies are integrated into response?
Analytic Monitoring: gather and analyze data on an ongoing basis and in a coordinated way to identify potential vulnerabilities, adversary activities, and damage	• **Sensor locations**: At how many locations is monitoring performed? • **Sensor coordination**: How well can sensor coverage and analysis be coordinated within the architecture? • **Sensor dynamism**: How quickly can sensors be recalibrated? • **Analysis timeliness**: How quickly can analysis of sensor or other data be performed? • **Scope**: What is the scope of analysis?
Coordinated Defense: manage adaptively and in a coordinated way multiple, distinct mechanisms to defend critical resources against adversary activities	• **Breadth of defense**: How many defensive techniques are applied at a given architectural layer? • **Depth of defense**: At how many architectural layers is a given defensive technique applied? • **Internal consistency / coordination**: How consistently and with how much coordination are cyber defenses, supporting security controls, and supporting performance controls managed within a given administrative span of control? • **External consistency / coordination**: How consistently and with how much coordination are cyber defenses, supporting security controls, and supporting performance controls managed across different administrative spans of control?

Cyber Resiliency Technique	Key Differentiators Between Levels
Deception: use obfuscation and misdirection (e.g., disinformation) to confuse an adversary	• **Sophistication of dissimulation**: How sophisticated are dissimulation / obfuscation techniques? • **Sophistication of simulation**: How sophisticated are simulation techniques? • **Integration**: How well are deception mechanisms coordinated / integrated with mechanisms from other practices?
Diversity: use a heterogeneous set of technologies (e.g., hardware, software, firmware, protocols) and data sources to minimize the impact of attacks and force adversaries to attack multiple different types of technologies	• **Depth of diversity**: At how many architectural layers is diversity provided or supported? • **Breadth of diversity**: At how many locations in the architecture is diversity provided or supported? • **Degree of diversity**: How many instances / alternatives are expected or accommodated within the selected architectural layers? • **Diversity dynamism**: How quickly (in terms of technology refreshes or response to incidents or vulnerability discoveries) can new implementations be integrated into the system? • **Integration**: How well is diversity integrated with other practices?
Dynamic Positioning: use distributed processing and dynamic relocation of critical assets and sensors	• **Asset positioning**: How extensively is a moving target defense strategy applied to critical assets? • **Sensor positioning**: How extensively can sensors be moved / reassigned / reconfigured in response to changes in the threat environment? • **Dynamism**: How quickly can dynamic positioning take effect?
Dynamic Representation: construct and maintain dynamic representations of components, systems, services, mission dependencies, adversary activities, and effects of alternative cyber courses of action	• **Breadth of representation**: How many aspects are included in representations? • **Timeliness of representation**: How quickly / how often are representations updated?
Non-Persistence: retain information, services, and connectivity for a limited time	• **Depth of non-persistence**: At how many architectural layers is non-persistence provided or supported? • **Frequency of non-persistence**: How frequently is the data, service, system, or connection refreshed?

Cyber Resiliency Technique	Key Differentiators Between Levels
Privilege Restriction: 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	• **Depth of privilege restriction**: At how many layers is privilege restriction applied? • **Breadth of privilege restriction**: How broadly or narrowly is least privilege applied (e.g., is it only applied to services, access to data, individuals)? • **Criticality**: To what degree is criticality analysis linked to least privilege? • **Coordination/consistency**: How consistently are privileges defined and assigned? In a system-of-systems, how well are policies and practices coordinated?
Realignment: align cyber resources with core aspects of mission/business functions, thus reducing the attack surface	• **Depth of realignment**: At how many layers is realignment applied? • **Degree of analysis**: How detailed is analysis/determination of core mission functions? • **Formalization of realignment**: How formal/structured are realignment processes (e.g., ad-hoc, ongoing, only in response to emergencies)?
Redundancy: maintain multiple protected instances of critical resources (information and services)	• **Breadth of redundancy**: How many duplicate copies of a given resource exist? Where? • **Depth of redundancy**: At how many layers is redundancy provided? • **Validation**: How consistent and independent are duplicate copies? • **Integration**: How well is redundancy integrated with other techniques? practices?
Segmentation: separate (logically or physically) components based on pedigree and/or criticality, to limit the spread of or damage from successful exploits	• **Strength of separation**: How effective is the separation? • **Depth of segmentation**: At how many layers is segmentation provided? • **Responsiveness of isolation**: How quickly and effectively can segmentation be used to isolate cyber resources in light of an attack?
Substantiated Integrity: ascertain that critical services, information stores, information streams, and components have not been corrupted by an adversary	• **Depth of integrity**: At how many layers is unpredictability applied? • **Strength of integrity mechanisms**: How strong or effective are the substantiated integrity mechanisms (e.g., prevent changes to data/system, detect changes, increase sources of data to reduce probability of changes that will impact mission)?
Unpredictability: make changes frequently and randomly, not just in response to actions by the adversary	• **Depth of unpredictability:** At how many layers is unpredictability applied? • **Intentionality of unpredictability:** Is unpredictability something that is planned, happenstance, or some combination?

5.2 Implementation

For notional architectures, the assessment is based on specifications and plans. Regardless of how detailed such documents are, they still focus on something that is not yet real. In contrast, "as-is" architectures are realized in operational environments. This means that there is more confidence in their assessments than in the assessments of the notional architectures. In addition, as-is architectures can be assessed with regards to the commitment to using a resiliency technique, the comprehensiveness of the implementation, and the effectiveness of the implementation. Note that implementation includes not only inclusion of technical mechanisms, but also how the practice is used operationally. Thus, the assessment takes into consideration CONOPS, SOPs, and CCoAs. Table 8 provides a general definition of levels of implementation; see Appendix C.3 for an example, using Diversity.

Table 8. Levels of Implementation

Level	Commitment	Comprehensiveness	Effectiveness
Very High	Policies and contractual agreements support the use of the technique Use of the technique is well integrated into operations (CONOPS, SOPs, TTPs) Resources are allocated to the use of the technique (life-cycle costs, LOE, training) Investment / architectural evolution plans include expected future mechanisms / capabilities	All specific technologies or approaches allowed by POET considerations are applied	Effectiveness validated by penetration testing, exercises, and metrics tracking
High	Policies and contractual agreements accommodate the use of the technique Use of the technique is integrated into operations (CONOPS, SOPs, TTPs) Resources are allocated to the use of the technique (life-cycle costs, LOE, training) Investment / architectural evolution plans include the technique	Most specific technologies or approaches allowed by POET considerations are applied	Effectiveness validated by penetration testing and limited exercises
Medium	Policies and contractual agreements accommodate some use of the technique Some uses of the technique are represented in operations (CONOPS, SOPs, TTPs) Limited resources are allocated to the use of the technique (life-cycle costs, LOE, training)	Some specific technologies or approaches allowed by POET considerations are applied	Effectiveness validated by testing
Low	Plans exist for modifying policies and contractual agreements to accommodate some use of the technique	Some specific technologies or approaches allowed by POET considerations are planned	Effectiveness to be validated by testing
Very Low	No plans to address POET considerations to enable or facilitate use of the technique	Techniques or approaches are incidental rather than planned	Effectiveness is not evaluated

In addition, an operational or "as-is" architecture can be assessed with respect to how well resilience-enhancing activities are performed. The general scoring model for resilience-enhancing activities uses a high/medium/low (or green/yellow/red) value scale, as shown in Table 9. The value of an assessment of resiliency-enhancing activities is that it tells a clear story of what can and cannot be done to achieve resiliency objectives. Table 22 in Appendix C.4 provides examples of value scales for activities.

Table 9. General Value Scale for Resilience-Related Cyber Defense Activities

Value	Description
High	The activity is performed in a rigorous and repeatable (or structured) manner, and is highly effective. Structured activities are well-documented, with the documentation including situation-dependent alternatives; are typically well supported by automation; and are informed by well-defined and well-maintained metrics. High corresponds roughly to Levels 4 and 5 of governance structures and processes in the Cyber Prep framework [7], or to Levels 4 and 5 in a capability maturity framework.
Medium	The activity is performed in a semi-structured manner, and is fairly effective. Semi-structured activities are documented and repeatable (with considerable variation), and are typically supported by some automation and/or metrics. Medium corresponds roughly to Level 3 of governance structures and processes in the Cyber Prep framework, or to Level 3 in a capability maturity framework.
Low	The activity is performed in an ad-hoc manner, with unpredictable effectiveness. Ad-hoc activities are typically reactive or after-the-fact, labor-intensive, undocumented and thus difficult to repeat consistently, and unsupported by metrics or automation. Low corresponds roughly to Levels 1 and 2 of governance structures and processes in the Cyber Prep framework, or to Levels 1 and 2 in a capability maturity framework.

6 Develop Recommendations

The goal of an assessment is to provide recommendations. This section provides general recommendations to serve as a starting point. Issues that architects and systems engineers should take into consideration when developing or applying recommendations are also discussed.

6.1 General Recommendations

As a starting point, a general set of recommendations has been derived from experience applying the cyber resiliency engineering framework, and from the Engineering Principles track at the Second Annual Secure and Resilient Cyber Architectures Workshop. Those recommendations are summarized in Tables 10 and 11.

Table 10 recommends cyber resiliency techniques to consider first. For these techniques, "low-hanging fruit" or opportunities for near-term and high-leverage improvements can be identified using Table 52 in Appendix F. The possible improvements identified from Table 52 can be tailored to the specific architecture, for example by identifying components, locations, or technologies to which recommendations apply. Table 11 provides general recommendations or engineering principles for applying cyber resiliency techniques.

Table 10. Recommendations for Assessment Priorities

Recommendations for Cyber Resiliency Techniques to Consider First	
• Use existing capabilities more effectively o Analytic Monitoring o Redundancy with Diversity o Privilege Restriction o Segmentation	• Address governance to enable existing and evolving capabilities to be more effective o Adaptive Response o Coordinated Defense o Analytic Monitoring

Table 11. General Recommendations for Applying Cyber Resiliency Techniques

Cyber Resiliency Technique	General Recommendations
Adaptive Response	• Maintain an up-to-date and consistent cyber playbook (set of SOPs, CCoAs, and configuration guides) o Exercise to validate • Integrate automated decision response mechanisms (e.g., dynamic reconfiguration) carefully, to avoid destabilization o Support human interaction and understandable user interfaces o Exercise caution in using fully automated dynamic mechanisms
Analytic Monitoring	• Combine monitoring and analysis across sub-systems (e.g., IDS, anti-malware, CMRS) • Identify and address monitoring issues related to transience of other cyber resources (e.g., end-user devices) • Analyze and address trade-off between encryption and monitoring
Coordinated Defense	• Apply defense in depth, moving away from a "hard outside, soft chewy center" • Coordinate the development of administrator SOPs, particularly for performance management and configuration / patch management, with mission threads • Coordinate the development of CCoAs with o Administrator SOPs, across multiple administrative domains o Mission threads, across missions that rely on resources covered by the CCoAs

Cyber Resiliency Technique	General Recommendations
Deception	• Work out policy, governance, and CONOPS issues related to active deception prior to defining a deception architecture • Consider the scope of deception (e.g., focused on internal systems, supply chain, DMZ, or external data repositories and servers) in architectural decisions
Diversity	• Make effective use of incidental diversity o Incorporate (rather than try to expunge) diverse components, products, and services acquired at different times and/or by different organizations o Accommodate diversity in end-user devices (particularly for BYOD) • Invest in targeted diversity for critical assets carefully o For communications, identify and maintain alternative communications paths o For software, take advantage of organization-owned mission applications apply implementation diversity o For information, identify and maintain multiple sources of critical mission data o For hardware, apply AT and SCRM as well as design diversity
Dynamic Positioning	• Use existing technologies to distribute assets in ways that take resiliency into account o Ensure consistent protection o Integrate with backup, isolation, and rollback
Dynamic Representation	• Ensure existence of and then build on static representations of components, systems, services mission dependencies and adversary actions • Use existing tools to maintain a current and realistic representation o Use Continuous Monitoring and intrusion detection tools to represent security posture o Use performance monitoring and map-the-mission tools to represent mission dependencies • Coordinate with contingency planning activities, so that plans, CCoAs, and SOPs can support non-adversarial as well as adversarial disruptions
Non-Persistence	• Leverage virtualization to make services non-persistent • Minimize “immortal” services and connections as part of system and network administrator SOPs o Terminate unused ports and protocols
Privilege Restriction	• Apply standards of good practice for least privilege, separation of duties, and role-based access control • Identify critical resources and lock down their use
Realignment	• Analyze mission/business processes to identify non-essential resources • Plan to separate or offload non-essential resources
Redundancy	• Apply standards of good practice for redundancy in the context of contingency planning • Ensure current patch/configuration status of redundant firmware and software resources • Ensure protection of all instances of critical resources regardless of location
Segmentation	• Define and separate enclaves based on sensitivity, criticality, and trust o Employ logical isolation mechanisms (e.g., routers, firewalls, controlled interfaces) to isolate enclaves and subnets o Ensure isolation of Internet from intranet o Isolate organization’s cyber security operations/response center from rest of organization.
Substantiated Integrity	• Apply existing software integrity and network address validation mechanisms effectively • Apply AT to critical hardware, firmware, and software components

Cyber Resiliency Technique	General Recommendations
Unpredictability	• Include unpredictable changes that are transparent to mission/business process users in day-to-day operations

When differentiating factors for cyber resiliency techniques are assessed, more specific recommendations can be developed. Examples for selected cyber resiliency techniques are presented in Appendix C. The recommendations in Appendix C can be tailored to provide guidance on how to instantiate or evolve the architecture in a way that maximizes the effective application of the selected cyber resiliency techniques.

6.2 Additional Considerations

The POET considerations identified in Appendix E should be considered when developing recommendations. A common theme in the set of technical considerations is the maturity of possible solutions. In addition, virtualization – both as enabling cyber resiliency solutions and as a challenge to some cyber resiliency techniques – should be considered. Finally, some considerations are specific to data centers.

6.2.1 Maturity of Resiliency Techniques

Not all of the resilience techniques are at that same level of maturity and usage. Some are more available or in common use today than others. Knowing this may be useful for those decision makers who are more prone to employ known commodities and find being early adopters disconcerting. Techniques such as Privilege Restriction, Redundancy, Segmentation, and Analytic Monitoring include elements that are in common practice today. At the other extreme, techniques such as Adaptive Response, Deception, Diversity, Realignment and Unpredictability are rarely applied in current practice.

These generalizations need to be tempered by considering the range of possible mechanisms the different cyber resiliency techniques include. For those techniques identified as in common practice today, it is often the case that only some elements are in common practice. Thus, for Analytic Monitoring it is generally the case that monitoring is in common practice, but other characteristics (e.g., finding indications of stealthy adversaries, detecting and assessing damage, watching for adversary activities during recovery and evolution) are not commonly done today. Similarly many organizations employ some elements of Diversity in their architectures (e.g., use of laptops, smartphones, and tablets). But in general it is an ad-hoc use of Diversity driven by factors such as cost and market leadership for a given domain. Very rarely is a diverse set of components employed systematically with the focus of making the adversary's work harder.

Finally, it is important to recognize that there are products and implementations to address all of the resiliency techniques in the near-, mid- and long-term timeframes. See Appendix F for more details.

6.2.2 Virtualization and Its Relationship to Resiliency

There is no single product or technology one can acquire that will automatically make an architecture more resilient. But the one technology that does seem to underlie effective implementation of resilience techniques more than any other is virtualization. Virtualization is an effective means of supporting Non-Persistence, allowing for the rapid refresh of services and data. Virtualization is also the means for supporting Diversity by allowing for the rapid deployment of diverse operating systems and applications. The rapid deployment capability of

virtualization means that it enables Adaptive Response [8]. Virtualization is also key to Dynamic Positioning. Virtualization can be used in conjunction with Segmentation strategies [8]. Virtualization's use and deployment is growing in modern architectures (e.g., cloud). But that does not meant that simply including virtualization technology in architecture makes the architecture more resilient. Most current uses of virtualization have little or no relevance to resiliency. Thus, cyber resiliency assessments should consider not just whether virtualization is in place, but also how effectively virtualization is used to provide resiliency.

6.2.3 Considerations for Data Centers

While the trend toward using data centers to host mission/business applications is largely driven by cost, data centers can provide increased resiliency, leveraging virtualization and cloud computing infrastructures. Cyber resiliency techniques apply differently to different classes of resources in data centers. Segmentation techniques to create separate enclaves for different classes of resources (and within some classes) are highly applicable to data center/cloud environments.

Housed resources are cyber resources owned and managed by another entity (organization or individual). The organization or facility provides a physical and/or networking infrastructure (e.g., power, HVAC, physical security; network connectivity), but has little or no insight into the operation of the resources. The organization or facility establishes minimum requirements for housed resources (e.g., range of power needs; security and/or interoperability requirements for network connectivity), but may have limited abilities to validate the entity's assertions that those requirements are met. Thus, Analytic Monitoring, Coordinated Defense, and Substantiated Integrity techniques may be difficult to apply, and Segmentation is important to protect house resources belonging to one entity from those belonging to another. In a data center/cloud environment, housed resources will typically be part of a Legacy Enclave.

Managed resources are cyber resources owned by another entity (organization or individual), which the organization or facility operates and manages on the entity's behalf. In addition to a physical and networking infrastructure, the organization or facility provides management and infrastructure services. In particular, the organization or facility provides security services and network operations. Coordinated Defense techniques thus become more feasible and important. In a data center/cloud environment, managed resources will typically be applications running in a cloud enclave. Segmentation is important to protect managed resources belonging to one entity from those belonging to another.

Owned and managed resources are cyber resources owned and managed by the organization or facility. In a data center/cloud environment, owned and managed resources are services and infrastructure which the organization or facility provides to managed resources (e.g., standard applications, backup and recovery services, security services), together with the services and infrastructure needed to provide those services (e.g., virtualization and cloud management services, core physical infrastructure). While all cyber resiliency techniques can be applied to owned and managed resources, the need to meet service level agreements must be considered, particularly for Adaptive Response techniques.

7 References/Bibliography

[1] M. Richard and J. Sain, *Cyber Kill Chain (PR 12-0886)*, Bedford, MA: The MITRE Corporation, 2012.

[2] The MITRE Corporation, "Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX)," 2012. [Online]. Available: http://measurablesecurity.mitre.org/docs/STIX-Whitepaper.pdf.

[3] D. Bodeau, R. Graubart, L. LaPadula, P. Kertzner, A. Rosenthal and J. Brennan, "Cyber Resiliency Metrics," April 2012. [Online]. Available: https://registerdev1.mitre.org/sr/12_2226.pdf.

[4] J. Allen and N. Davis, "Measuring Operational Resilience Using the CERT® Resilience Management Model," September 2010. [Online]. Available: http://www.cert.org/archive/pdf/10tn030.pdf.

[5] DOE, "Electricity Subsector Cybersecurity Capability Maturity Model (ES-C2M2)," 31 May 2012. [Online]. Available: http://energy.gov/sites/prod/files/Electricity%20Subsector%20Cybersecurity%20Capabilities%20Maturity%20Model%20(ES-C2M2)%20-%20May%202012.pdf.

[6] D. Bodeau and R. Graubart, "Cyber Resiliency Engineering Framework, Version 1.0," September 2011. [Online]. Available: http://www.mitre.org/work/tech_papers/2012/11_4436/11_4436.pdf.

[7] D. J. Bodeau, R. D. Graubart and J. Fabius-Greene, "Cyber Security Governance," September 2010. [Online]. Available: http://www.mitre.org/work/tech_papers/2010/10_3710/10_3710.pdf.

[8] J. M. Butler and R. Vandenbrink, "IT Audit for the Virtual Environment," September 2009. [Online]. Available: http://www.sans.org/reading_room/analysts_program/VMware_ITAudit_Sep09.pdf.

[9] J. Chan, "Virtualization: IT Audit and Security Perspectives (presentation)," March 2011. [Online]. Available: http://www.universityofcalifornia.edu/compaudit/documents/symposium2010/Chan.pdf.

[10] R. Pietravalle and D. Lanz, "Resiliency Research Snapshot (PR Case No. 11-3023)," June 2011. [Online]. Available: http://www.mitre.org/work/tech_papers/2011/11_3023/11_3023.pdf.

[11] Fedora Project, "Dynamic Firewall with firewalld," 21 April 2011. [Online]. Available: http://fedoraproject.org/wiki/Features/DynamicFirewall.

[12] E. Ipek, M. Kirman, N. Kirman and J. F. Martinez, "Core Fusion: Accommodating Software Diversity in Chip Multiprocessors," in *Proceedings of the 34th International Symposium on Computer Architecture*, San Diego, CA, 2007.

[13] P. Cotret, G. Gogniat, J.-P. Diguet and J. Crenne, "Lightweight Reconfiguration Security Services for AXI-Based MPSoCs," in *Proceedings of the 22nd International Conference on Field Programmable Logic and Applications (FPL 2012)*, Oslo, Norway, 2012.

[14] C. T. Rathgeb and G. D. Peterson, "Secure Processing Using Dynamic Partial Reconfiguration," in *Proceedings of the 5th Annual Cyber Security and Information Intelligence Research Workshop (CSIIRW '09)*, Knoxville, TN, 2009.

[15] V. Zaborovsky, V. Mulukha, A. Silinenko and S. Kupreenko, "Dynamic Firewall Configuration: Security System Architecture and Algebra of the Filtering Rules," 19-24 June 2011. [Online]. Available: http://www.thinkmind.org/download.php?articleid=internet_2011_2_20_30062.

[16] M. Tkačik, K. Kleinová and P. Feciľak, "Dynamic network reconfiguration based on application measurements with the goal of network traffic optimization," in *Proceedings of 9th International Conference on Emerging eLearning Technologies and Applications (ICETA)*, Stará Lesná, The High Tatras, Slovakia, 2011.

[17] G. Dan Mois, S. Flonta, S. Iulia, S. Enyedi and L. C. Miclea, "Reconfiguration security for hardware agents in testing," in *Proceedings of the 2010 International Conference on Automation Quality and Testing Robotics (AQTR)*, Cluj-Napoca, Romania , 2010.

[18] L. Tan and A. Krings, "An Adaptive N-variant Software Architecture for Multi-Core Architectures: Models and Performance Analysis," in *Proceedings of the 2011 International Conference on Computational Science and its Applications (ICCSA'11)*, Santander, Spain, 2011.

[19] D. Cañas, M. B. Crouse and E. W. Fulp, "Improving the Diversity Defense of Genetic Algorithm-Based Moving Target Approaches," 11 June 2012. [Online]. Available: http://cps-vo.org/file/3707/download/9982.

[20] R. Rodrigues, B. Liskov, K. Chen, M. Liskov and D. Schultz, "Automatic Reconfiguration for Large-Scale Reliable Storage Systems," *IEEE Transactions on Secure and Dependable Computing*, pp. 145-158, March-April 2012.

[21] A. Chandra, W. Gong and P. Shenoy, "Dynamic Resource Allocation for Shared Data Centers Using Online Measurements," *ACM SIGMETRICS Performance Evaluation Review*, pp. 300-301, June 2003.

[22] G. Soundararajan, D. Lupei, S. Ghanbari, A. D. Popescu, J. Chen and C. Amaza, "Dynamic Resource Allocation for Database Servers Running on Virtual Storage," in *Proceedings of 7th USENIX Conference on File and Storage Technologies*, San Francisco, CA, 2009.

[23] V. F. Farias and B. Van Roy, "Approximation Algorithms for Dynamic Resource Allocation," 26 December 2004. [Online]. Available: http://www.stanford.edu/~bvr/psfiles/project.pdf.

[24] J. Guitart, D. Carrera, V. Beltran, J. Torres and E. Ayguade, "Dynamic CPU provisioning for self-managed secure web applications in SMP hosting platforms," *Computer Networks: The International Journal of Computer and Telecommunications Networking*, vol. 52, no. 7, pp. 1390-1409, 2008.

[25] Ciena, "Ciena's Dynamic Resource Allocation Controller Released as Open Source Software," 30 June 2010. [Online]. Available: http://www.ciena.com/corporate/news-events/press-releases/Cienas-Dynamic-Resource-Allocation-Controller-Released-As-Open-Source-Software-.html.

[26] S. Song and K. Hwang, "Dynamic Grid Security with Trust Integration and Dynamic Resource Allocation," 2004. [Online]. Available: http://gridsec.usc.edu/files/publications/SongHPDC.pdf.

[27] T. Graves, T. Reale and C. Schmidt, "Dynamic Composability (PR 11-1575)," 2011. [Online]. Available: http://www.mitre.org/work/areas/research/2011iebriefings/05MSR160-JP.pdf.

[28] S. Hariharasubrahmanian, "Dynamic composability building flexible complex real-time systems," 1 January 2003. [Online]. Available: http://scholarworks.umass.edu/dissertations/AAI3079596/.

[29] R. Want, T. Pering, S. Sud and B. Rosario, "Dynamic Composable Computing (DCC)," 2008. [Online]. Available: http://berkeley.intel-research.net/~brosario/papers/Hotmobile08%20DCC.pdf.

[30] M. Castro and B. Liskov, "Practical Byzantine fault tolerance and proactive recovery," *ACM Transactions on Computer Systems*, pp. 398-461, 2002.

[31] A. F. Hansen, O. Lysne, T. Cicik and S. Gjessing, "Fast Proactive Recovery from Concurrent Failures," in *Proceedings of the IEEE International Conference on Communications (ICC '07)*, Glasgow, Scotland, 2007.

[32] P. Sousa, N. F. Neves and P. Verissimo, "Hidden Problems of Asynchronous Proactive Recovery," in *Proceedings of the 3rd workshop on on Hot Topics in System Dependability (HotDep'07)*, Edinburgh, UK, 2007.

[33] P. Sousa, N. F. Neves, P. Verissimo and W. H. Sanders, "Proactive Resiliency Revisited: The Delicate Balance Between Resisting Intrusions and Remaining Available," in *Proceedings of the 25th IEEE Symposium on Reliable Distributed Systems (SRDS'06)*, Leeds, UK, 2006.

[34] T. Distler, R. Kapitza and H. P. Reiser, "State Transfer for Hypervisor-Based Proactive Recovery of Heterogeneous Replicated Services," 2010. [Online]. Available: http://www4.cs.fau.de/~distler/publications/distler10state.pdf.

[35] F. Zhao, M. Li, W. Qiang, H. Jin, D. Zou and Q. Zhang, "Proactive recovery approach for intrusion tolerance with dynamic configuration of phsyical and virtual replicas," *Security and Communication Networks*, pp. 1169-1180, October 2012.

[36] P. Barford, M. Dacier, T. G. Dietterich, M. Fredrikson, J. Griffin, S. Jajodia, S. Jha, J. Li, P. Liu, P. Ning, X. Ou, D. Song, L. Strater, V. Swarup, G. Tadda, C. Wang and J. Yen, "Cyber SA: Situational Awareness for Cyber Defense," in *Cyber Situational Awareness: Issues and Research (Advances in Information Security, Vol. 46)*, Springer, 2010, pp. 3-35.

[37] The MITRE Corporation, "The National Security Engineering Center," July 2012. [Online]. Available: http://www.mitre.org/news/pdfs/nsec.pdf.

[38] I. Gul, A. ur Rehman and M. H. Islam, "Cloud Computing Security Auditing," in *Proceedings of the 2nd International Conference on Next Generation Information Technology*, Gyeongju City, Republic of Korea, 2011.

[39] M. Sharif, W. Lee and W. Cui, "Secure In-VM Monitoring Using Hardware Virtualization," in *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS '09)*, Chicago, IL, 2009.

[40] J. Spring, "Monitoring Cloud Computing by Layer, Part 1," *IEEE Security & Privacy*, March/April 2011.

[41] J. Spring, "Monitoring Cloud Computing by Layer, Part 2," *IEEE Security & Privacy*, pp. 52-55, May-June 2011.

[42] NIST, "Guide to Security for Full Virtualization Technologies, NIST SP 800-125," January 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-125/SP800-125-final.pdf.

[43] A. V. Taddeo, L. G. G. Morales and A. Ferrante, "System Policies for Gradual Tuning of

Security and Workload in Wireless Sensor Networks," in *Proceedings of the 2011 Wireless Telecommunications Symposium (WTS)*, New York, NY, 2011.

[44] S. Xu, "Toward a theoretical framework for trustworthy cyber sensing," *Proc. SPIE 7709, Cyber Security, Situation Management, and Impact Assessment II; and Visual Analytics for Homeland Defense and Security II*, *770907*, 28 April 2010.

[45] S. Mandala, M. A. Ngadi and A. H. Abdullah, "A Survey on MANET Intrusion Detection," *International Journal of Computer Science and Security,* vol. 2, no. 1, 2008.

[46] D. T. Stott, L. G. Greenwald, O. P. Kreidl and B. DeCleene, "Tolerating Adversaries in the Estimation of Network Parameters from Noisy Data: A Nonlinear Filtering Approach," in *Proceedings of MILCOM 2009*, Boston, MA, 2009.

[47] M. Sunu, S. Upadhyaya, M. Sudit and A. Stotz, "Situation Awareness of Multistage Cyber Attacks by Semantic Data Fusion," in *Proceedings of the 2010 Military Communications Conference (MILCOM'10)*, San Jose, CA, 2010.

[48] Bivio Networks, "NetFalcon: The First Big Data Solution for Enterprise Network Traffic," 2012. [Online]. Available: http://www.bivio.net/public_pdfs/Bivio_DS_NetFalcon_ENT.pdf.

[49] Pervasive, "Pervasive DataRush: A New Tool to Help Win the Cyber Security War," 2011. [Online]. Available: http://bigdata.pervasive.com/wdata/pp/document/download/06930000001Ta96AAC.

[50] P. Frühwirt, M. Huber, M. Mulazzani and E. R. Weippl, "InnoDB Database Forensics," April 2010. [Online]. Available: http://www.sba-research.org/wp-content/uploads/publications/AINA2010-InnoDBforensics_preprint.pdf.

[51] M. Olivier, "On metadata context in Database Forensics," *Digital Investigation, Vol. 5,* pp. 115-123, March 2009.

[52] J. Scholtz and A. Narayanan, "Towards an Automated Digital Data Forensic Model with specific reference to Investigation Processes," in *Proceedings of the 8th Australian Digital Forensics Conference*, Perth, Australia, 2010.

[53] S. Musman, A. Temin, M. Tanner, D. Fox and B. Pridemore, "Evaluating the Impact of Cyber Attacks on Missions (PR Case No. 09-4577)," 2009. [Online]. Available: http://198.49.146.10/work/tech_papers/2010/09_4577/09_4577.pdf.

[54] P. Liu, J. Jia, S. Zhang, Y.-C. Jhi, K. Bai and J. Li, "Cross-Layer Damage Assessment for Cyber Situational Awareness," in *Advances in Information Security, 2010, Volume 46, Part 4*, Springer, 2010, pp. 155-176.

[55] ISF, "Data Analytics for Information Security: From hindsight to insight," 1 August 2012. [Online]. Available: https://www.securityforum.org/userfiles/public/isf_data_analytics_executive-summary.pdf.

[56] K. A. McVearry, "Pedigree Management and Assessment Framework (PMAF) Demonstration," in *Proceedings of the 3rd International Provenance and Annotation Workshop (IPAW)*, Troy, NY, 2010.

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

[58] CA Technologies, "Dynamic Cyber Defense," 2010. [Online]. Available: http://www.ca.com/~/media/Files/whitepapers/ca-dynamic-cyberdefense.pdf.

[59] T. Bandyopadhyay, "A Model for B2B IT Security: Multilayer Defense Facing Interdependent Cyber Risk," in *Proceedings of the Sixteenth Annual Conference of the Southern Association for Information Systems (SAIS 2011)*, Atlanta, GA, 2011.

[60] M. H. Shuyuan, "A Framework of Coordinated Defense," in *Second International Conference on Computational Cultural Dynamics*, University of Maryland College Park, MD, 2008.

[61] S. Foote, M. Kramer and B. Yost, "Alternative Processes and Operations Controlled via a Cyber Operations Center (CyOC), PR # 11-1567," 2011. [Online]. Available: http://www.mitre.org/work/areas/research/2011iebriefings/05MSR160-JO.pdf.

[62] D. A. Rodriguez-Silva, F. J. Gonzalez-Castano, L. Adkinson-Orellana, A. Fernandez-Cordeiro, J. R. Troncoso-Pastoriza and D. Gonzalez-Martinez, "Encrypted Domain Processing for Cloud Privacy," in *Proceedings of the 1st International Conference on Cloud Computing and Services (CLOSER '11)*, Noordwijkerhout, The Netherlands, 2011.

[63] R. S. Chakraborty, "Hardware Security Through Design Obfuscation," May 2010. [Online]. Available: http://etd.ohiolink.edu/send-pdf.cgi/Chakraborty%20Rajat%20Subhra.pdf?case1270133481.

[64] Allure Security Technology, Inc., "Final Report: Anomaly Detection At Multiple Scales (ADAMS)," 11 November 2011. [Online]. Available: http://www.dtic.mil/dtic/tr/fulltext/u2/a552461.pdf.

[65] D. ". Ragsdale, "Scalable Cyber Deception," in *DARPA Cyber Colloquium*, Arlington, VA, 2011.

[66] ReSIST, "Resilience-Explicit Computing: final," December 2008. [Online]. Available: http://www.resist-noe.org/Publications/Deliverables/D33-ResEx_computing_final.pdf.

[67] P. Trimintzios, C. Hall, R. Clayton, R. Anderson and E. Ouzounis, "Inter-X: Resilience of the Internet Interconnection Ecosystem," April 2011. [Online]. Available: http://www.enisa.europa.eu/act/res/other-areas/inter-x/report/interx-report/at_download/fullReport.

[68] M. G. Richards, A. M. Ross, D. E. Hastings and D. H. Rhodes, "Empirical Validation of Design Principles for Survivable System Architecture," in *Proceedings of the 2nd Annual IEEE Systems Conference*, Montreal, Quebec, Canada, 2008.

[69] H. Goldman, "Building Secure, Resilient Architectures for Cyber Mission Assurance," 2010. [Online]. Available: http://www.mitre.org/work/tech_papers/2010/10_3301/10_3301.pdf.

[70] B. A. Allan, R. C. Armstrong, J. R. Mayo, L. G. Pierson, M. D. Torgerson and A. M. Walker, "The Theory of Diversity and Redundancy in Information System Security: LDRD Final Report," October 2010. [Online]. Available: http://prod.sandia.gov/techlib/access-control.cgi/2010/107055.pdf.

[71] J. Han, D. Gao and R. H. Deng, "On the Effectiveness of Software Diversity: A Systematic Study on Real-World Vulnerabilities," July 2009. [Online]. Available: http://www.mysmu.edu/phdis2007/jin.han.2007/DIMVA09Han.pdf.

[72] K. Dulaney, "Use Managed Diversity to Support the Growing Variety of Endpoint Devices," Gartner, 2011.

[73] J. Schekkerman, "STREAM: A Successful and Pragmatic "Managed Diversity‟ Enterprise Architecture Approach," 2012. [Online]. Available: http://www.enterprise-architecture.info/Images/STREAM/White%20Paper%20STREAM%20%202010%20v1.2

.pdf.

[74] M. J. van der Meulen and M. Revilla, "The Effectiveness of Software Diversity in a Large Population of Programs," *IEEE Transactions on Software Engineering, Vol. 34, No. 6,* pp. 753-764, November/December 2008.

[75] A. Sharma, "Reviewing Hardware Design Diversity," 14 July 2012. [Online]. Available: http://www.ee.ucla.edu/~ankur/parent/research/designDiversity/reviewing_design_diversity_Jul21.pdf.

[76] University of Bristol, "Enhancing software diversity project," 28 November 2007. [Online]. Available: http://www.bristol.ac.uk/spc/research/diversity.html.

[77] H. Wang and N. Liang, "A Software Diversity Model for Embedded Safety-critical System," in *Proceedings of 2009 Conference on Wireless Networks and Information Systems (WNIS'09)*, Shanghai, China, 2009.

[78] P.-y. Chen, G. Kataria and R. Krishnan, "Correlated Failures, Diversification and Information Security Risk Management," *MIS Quarterly*, pp. 397-422, June 2011.

[79] J. Boyens, C. Paulsen, N. Bartol, R. Moorthy and S. Shankles, "Notional Supply Chain Risk Management Practices for Federal Information Systems," March 2012. [Online]. Available: http://csrc.nist.gov/publications/drafts/nistir-7622/second-public-draft_nistir-7622.pdf.

[80] J. E. Just and M. Cornwell, "Review and analysis of synthetic diversity for breaking monocultures," in *Proceedings of the 2nd Workshop on Rapid Malcode (WORM'04)*, Fairfax, VA, 2004.

[81] A. Gherbi, R. Charpentier and M. Couture, "Software Diversity for Future Systems Security," *CrossTalk,* vol. 24, no. 5, pp. 10-13, September/October 2011.

[82] A. Vignaux, A. Auguste, B. Korel, S. Ren and K. Kwiat, "Improving Operation Time Bounded Mission Critical Systems' Attack-Survivability through Controlled Source-Code Transformation," July 2011. [Online]. Available: http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA560899.

[83] B. Cox, D. Evans, A. Filipi, J. Rowanhill, W. Hu, J. Davidson, J. Knight, A. Nguyen-Tuong and J. Hiser, "N-Variant Systems: A Secretless Framework for Security through Diversity," in *Proceedings of the 2006 USENIX Annual Technical Conference (USENIX '06)*, Boston, MA, 2006.

[84] M. Franz, "E unibus pluram: Massive-Scale Software Diversity as a Defense Mechanism," in *Proceedings of the 2010 New Security Paradigms Workshop*, Concord, MA, 2010.

[85] D. Williams, W. Hu, J. W. Davidson, J. D. Hiser, J. C. Knight and A. Nguyen-Tuong, "Security Through Diversity: Leveraging Virtual Machine Technology," *IEEE Security & Privacy*, January/February 2009.

[86] J. Rowe, K. N. Levitt, T. Demir and R. Erbacher, "Artificial Diversity as Maneuvers in a Control Theoretic Moving Target Defense," 11 June 2012. [Online]. Available: http://cps-vo.org/file/3710/download/9981.

[87] R. Albuquerque, J. Casper, E. Cheung, R. Couture, B. Lai, P. Leveille, J. Hu and D. Mauer, "Improving Data Analysis Through Diverse Data Source Integration," in *MILCOM*, Boston, MA, 2009.

[88] B. Glavic and K. Dittrich, "Data Provenance: A Categorization of Existing Approaches," in *Proceedings of the Datenbanksysteme für Business, Technologie und Web (BTW) Conference*, Aachen, Germany, 2007.

[89] A. Nguyen-Tuong, D. Evans, J. C. Knight, B. Cox and J. W. Davidson, "Security through Redundant Data Diversity," in *38th IEEE/IFPF International Conference on Dependable Systems and Networks, Dependable Computing and Communications Symposium*, Anchorage, AK, 2008.

[90] S. Forrest and E. G. Barrantes, "Increasing Communications Security through Protocol Parameter Diversity," August 2006. [Online]. Available: http://crypto.stanford.edu/portia/papers/paramRAN-CLEI-2006.pdf.

[91] D. Tse and P. Viswanath, "Point-to-point communication: detection, diversity, and channel uncertainty," in *Fundamentals of Wireless Communication*, Cambridge University Press, 2005, pp. 49-71.

[92] Cyber Physical Systems Virtual Organization, "National Symposium on Moving Target Research," 11 June 2012. [Online]. Available: http://cps-vo.org/group/mtrs/program.

[93] S. Jajodia, A. K. Ghosh, V. Swarup, C. Wang and X. S. Wang, Moving Target Defense: Creating Asymmetric Uncertainty for Cyber Threats (Advances in Information Security, Vol. 54), Springer, 2011.

[94] P. Beraud, A. Cruz, S. Hassell and S. Meadows, "Using Cyber Maneuver to Improve Network Resiliency," in *MILCOM*, Baltimore, MD, 2011.

[95] B. Parno, J. R. Lorch, J. R. Douceur, J. Mickens and J. M. McCune, "Memoir: Practical State Continuity for Protected Modules," in *Proceedings of the 2011 IEEE Symposium on Security and Privacy*, Oakland, CA, 2011.

[96] H. Okhravi, E. I. Robinson, A. Cornella, S. Yannalfo, P. W. Michaleas and J. Haines, "TALENT: Dynamic Platform Heterogeneity for Cyber Survivability of Mission Critical Applications," 29 October 2010. [Online]. Available: http://www.ll.mit.edu/mission/communications/publications/publication-files/full_papers/2010_10_29_Okhravi_SRCA_FP.pdf.

[97] J. H. Jafarian, E. Al-Shaer and Q. Duan, "OpenFlow Random Host Mutation: Transparent Moving Target Defense using Software Defined Networking," in *Proceedings of ACM SIGCOMM 2012*, Helsinki, Finland, 2012.

[98] F. Li, "Create Moving Target Defense in Static Networks by Learning from Botnets," in *Proceedings of the 13th Annual CERIAS Information Security Symposium*, West Lafayette, IN, 2012.

[99] S. Kuhn and S. Taylor, "Increasing Attacker Workload with Virtual Machines," in *MILCOM*, Baltimore, MD, 2011.

[100] M. Dunlop, S. Groat, W. Urbanski, R. C. Marchany and J. G. Tront, "MT6D: A Moving Target IPv6 Defense," in *MILCOM*, Baltimore, MD, 2011.

[101] W. Peters, "Mission Based Analysis for Cyber Measurement and Mission Assurance," in *Proceedings of the 5th Annual IT Security Automation Conference*, Baltimore, MD, 2009.

[102] F. C. Belz, "Space Segment Information Assurance Guidance for Mission Success," Aerospace Report No. TOR-2011(8591)-22, 2011.

[103] G. Jakobson, "Mission Cyber Security Situation Assessment Using Impact Dependency Graphs," in *Proceedings of the 14th International Conference on Information Fusion*, Chicago, IL, 2011.

[104] R. Layland, "The dark side of server virtualization," 7 July 2010. [Online]. Available: http://www.networkworld.com/news/2010/070610-virtualization-dark-side.html.

[105] K. Hildrum, F. Douglis, J. L. Wolf and P. S. Yu, "Storage Optimization for Large-Scale Distributed Stream-Processing Systems," *ACM Transactions on Storage, Vol. 3, No. 4*, 2008.

[106] R. Durst, S. Jajodia, A. Mei and S. Symington, "A Secure, Structured, Distributed Caching System for Providing Availability of Mission-Critical Reference Data," 2009. [Online]. Available: http://www.mitre.org/work/tech_papers/tech_papers_09/09_2115/09_2115.pdf.

[107] Y. Tang, P. P. Lee, J. C. Lui and R. Perlman, "FADE: Secure Overlay Cloud Storage with File Assured Deletion," in *Proceedings of the 6th Iternational ICST Conference, SecureComm 2010*, Singapore, 2010.

[108] P. Stahlberg, G. Miklau and B. N. Levine, "Threats to Privacy in the Forensic Analysis of Database Systems," in *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, Beijing, China, 2007.

[109] M. Wei, L. M. Grupp, F. E. Spada and S. Swanson, "Reliably Erasing Data From Flash-Based Solid State Drives," February 2011. [Online]. Available: http://static.usenix.org/event/fast11/tech/full_papers/Wei.pdf.

[110] NIST and NSA, "NISTIR 7657, A Report on the Privilege (Access) Management Workshop," March 2010. [Online]. Available: http://csrc.nist.gov/publications/nistir/ir7657/nistir-7657.pdf.

[111] J. Shin, V. Zyuban, P. Bose and T. M. Pinkston, "A Proactive Wearout Recovery Approach for Exploiting Microarchitectural Redundancy to Extend Cache SRAM Lifetime," in *Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA '08)*, Beijing, China, 2008.

[112] Invensys, "Mesh Network," 2012. [Online]. Available: http://iom.invensys.com/EN/Pages/Foxboro_DCSIASeries_MeshNetwork.aspx.

[113] J. Brand and R. Bach, "Overcoming Stressed Satellite Networks Using Alternative Communications," in *MILCOM*, Boston, MA, 2009.

[114] Cisco, "Market Data Network Architecture (MDNA) Overview," 2008. [Online]. Available: http://www.cisco.com/web/strategy/docs/finance/md-arch-ext.pdf.

[115] R. Durst and S. Symington, "Data Integrity and Availability: Using Mission Assurance Through Availability (MATA) - Resilient Assured Multiprovider Storage (RAMS)," 2011. [Online]. Available: http://www.mitre.org/work/areas/research/2011iebriefings/05MSR160-JA.pdf.

[116] A. W. Kosner, "Survey of Effects of Cloud Outage Shows How Much of the Web Runs on Amazon," 1 July 2012. [Online]. Available: http://www.forbes.com/sites/anthonykosner/2012/07/01/survey-of-effects-of-cloud-outage-shows-how-much-of-the-web-runs-on-amazon/.

[117] Palo Alto Networks, "Redundancy & Resiliency," 2012. [Online]. Available: http://www.paloaltonetworks.com/products/features/redundancy-resiliency.html.

[118] H. Cam, "RISE: Relational-Integrity-Sensitive-Encoding for Program Intrusion Detection," in *Cyber Sensing 2012, SPIE Defense, Security, and Sensing*, Baltimore, MD, 2012.

[119] Cimcor, "Advanced File Integrity Monitoring for IT Secuity, Integirty, and Compliance," 2011. [Online]. Available: http://www.cimcor.com/lp/assets/CimTrak_FIM_Whitepaper.pdf.

[120] TripWire, "File Integrity Monitoring," 2010. [Online]. Available: http://www.tripwire.com/asset/file-integrity-monitoring-secure-your-virtual-and-physical-it-environments/.

[121] A. Vasudevan, J. M. McCune, N. Qu, L. van Doorn and A. Perrig, "Requirements for an Integrity-Protected Hypervisor on the x86 Hardware Virtualized Architecture," Trust and Trustworthy Computing (Trust 2010), June 2010. [Online]. Available: http://users.ece.cmu.edu/~jmmccune/papers/vasudevan_mccune_ning_leendert_perrig_sechyp_trust2010.pdf.

[122] H. Langweg, "Building a Trusted Path for Applications Using COTS Components," in *RTO IST Symposium on "Adaptive Defence in Unclassified Networks" (RTO-MP-IST-041)*, Toulouse, France, 2004.

[123] Z. Zhou, V. D. Gligor, J. Newsome and J. M. McCune, "Building Verifiable Trusted Path on Commodity x86 Computers," IEEE Symposium on Security and Privacy, May 2012. [Online]. Available: http://users.ece.cmu.edu/~jmmccune/papers/ZhGlNeMc2012.pdf.

[124] D. Martinenghi, "Advanced Techniques for Efficient Data Integrity Checking," 2005. [Online]. Available: http://akira.ruc.dk/~henning/publications/MC_ADBIS2005.pdf.

[125] HDIV, "HTTP Data Integrity Validator," January 2012. [Online]. Available: http://www.hdiv.org/.

[126] The MITRE Corporation, "The TPM and OVAL: Using the Trusted Platform Module to Enhance OVAL Assessments," 2011. [Online]. Available: http://oval.mitre.org/language/about/docs/OVAL_and_TPM_White_Paper.pdf.

[127] Y. Li, J. M. McCune and A. Perrig, "VIPER: Verifying the Integrity of PERipherals' Firmware," in *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS)*, Chicago, IL, 2011.

[128] NIST, "BIOS Protection Guidelines, NIST SP 800-147," April 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-147/NIST-SP800-147-April2011.pdf.

[129] W3C Incubator Group, "Provenance XG Final Report," 8 December 2010. [Online]. Available: http://www.w3.org/2005/Incubator/prov/XGR-prov-20101214/#A_Working_Definition_of_Provenance.

[130] E. Bertino, C. Dai, H.-S. Lim and D. Lin, "High-Assurance Integrity Techniques for Databases," in *Proceedings of the 25th British National Conference on Databases: Sharing Data, Information and Knowledge (BNCOD '08)*, Cardiff, UK, 2008.

[131] A. Moitra, B. Barnett, A. Crapo and S. J. Dill, "Data Provenance Architecture to Support Information Assurance in a Multi-Level Secure Environment," in *MILCOM*, Boston, MA, 2009.

[132] A. Moitra, B. Barnett, A. Crapo and S. J. Dill, "Addressing Uncertainty and Conflicts in Cross-Domain Data Provenance," in *MILCOM*, San Jose, CA, 2010.

[133] M. Kasunic, D. Zubrow and E. Harper, "Issues and Opportunities for Improving the Quality and Use of Data in the Department of Defense," March 2011. [Online]. Available: http://www.sei.cmu.edu/reports/11sr004.pdf.

[134] GRDI, "Global Research Data Infrastructure (GRDI) 2020: Data Provenance and Trust," March 2011. [Online]. Available: http://grdi2020.eu/mediawiki/index.php/Data_Provenance_and_Trust.

[135] J. Park, D. Nguyen and R. Sandhu, "A Provenance-based Access Control Model," in *Proceedings of the 10th IEEE Conference on Privacy, Security and Trust (PST)*, Paris,

France, 2012.

[136] A. N. Bessani, "From Byzantine Fault Tolerance to Intrusion Tolerance," in *Proceedings of the 5th 5th Workshop on Recent Advances in Intrusion-Tolerant Systems (WRAITS 2011)*, Hong Kong, China, 2011.

[137] J. Kirsch, "Intrusion-Tolerant Replication under Attack," 2010. [Online]. Available: http://www.dsn.jhu.edu/~yairamir/JonKirsch_thesis.pdf.

[138] H. P. Reiser, "Byzantine Fault Tolerance for the Cloud," 2011. [Online]. Available: http://www.zurich.ibm.com/~cca/csc2011/submissions/reiser.pdf.

[139] A. Bessani, M. Correia, B. Quaresma, F. Andre and P. Sousa, "DEPSKY: Dependable and Secure Storage in a Cloud-of-Clouds," in *EuroSys '11: Proceedings of the Sixth Conference on Computer Systems*, Salzburg, Austria, 2011.

[140] O. Whitehouse, "An Analysis of Address Space Layout Randomization on Windows Vista," 2007. [Online]. Available: http://www.symantec.com/avcenter/reference/Address_Space_Layout_Randomization.pdf

[141] J. Kruse, S. Landsman, P. Smyton, A. Dziewulski, H. Hawley and M. King, "The POET Approach: A collaborative means for enhancing C2 systems engineering," in *Proceedings of the International Command and Control Research and Technology Symposium*, Fairfax, VA, 2012.

[142] T.-P. Liang, J. Jiang, G. S. Klein and J. Y.-C. Liu, "Software Quality as Influenced by Informational Diversity, Task Conflict, and Learning in Project Teams," *IEEE Transactions on Engineering Management, Vol. 57, No. 3*, pp. 477-486, August 2012.

[143] J. Bell and B. Whaley, Cheating and Deception, New Brunswick, NJ: Transaction Publishers, 1982.

[144] DoD, "Defense Acquisition Guidebook," 9 July 2011. [Online]. Available: http://www.dote.osd.mil/docs/dote-temp-guidebook/DEFENSE-ACQUISITION-GUIDEBOOK-07-29-2011.pdf.

[145] DoD, "Systems Engineering Guide for Systems of Systems, Version 1.0," August 2008. [Online]. Available: http://www.acq.osd.mil/se/docs/SE-Guide-for-SoS.pdf.

[146] Oracle, "An Oracle White Paper in Enterprise Architecture—Information Architecture: An Architect's Guide to Big Data," August 2012. [Online]. Available: http://www.oracle.com/technetwork/topics/entarch/articles/oea-big-data-guide-1522052.pdf.

[147] IBM, "Crisis simulation exercise," 2006. [Online]. Available: http://www-935.ibm.com/services/uk/its/pdf/crisis_management_06_its_001943.pdf.

[148] ASD(CIIA), "Deputy Assistant Secretary of Defense for Cyber, Identity, and Information Assurance Strategy," August 2009. [Online]. Available: http://cio-nii.defense.gov/docs/DoD_IA_Strategic_Plan.pdf.

[149] Civolution, "Digital Fingerprinting: White Paper," 2010. [Online]. Available: http://www.mediahedge.com/fileadmin/bestanden/pdf/White_Paper_-_Digital_Fingerprinting_by_Mediahedge_01-2010.pdf.

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

[151] OMB, "Improving Agency Performance Using Information and Information Technology,"

June 2009. [Online]. Available: http://www.whitehouse.gov/sites/default/files/omb/assets/fea_docs/OMB_EA_Assessment_Framework_v3_1_June_2009.pdf.

[152] DoD, "Interoperability and Supportability of Information Technology (IT) and National Security Systems (NSS)," 23 April 2007. [Online]. Available: http://www.dtic.mil/whs/directives/corres/pdf/463005p.pdf.

[153] Trail of Bits, "iVerify: iOS Integrity Verification," [Online]. Available: http://www.trailofbits.com/products/.

[154] M. Westphal, "JCIDS and Architectures: An Update," March 2012. [Online]. Available: http://www.dodenterprisearchitecture.org/program/Documents/WestphalJCIDS%20Update%202%20Mar2012%20FINAL%20[Compatibility%20Mode].pdf.

[155] L. S. Tinnel, O. S. Saydjari and D. Farrell, "Cyberwar Strategy and Tactics: An Analysis of Cyber Goals, Strategies, Tactics, and Techniques," in *Proceedings of the 2002 IEEE Workshop on Information Assurance*, West Point, NY, 2002.

[156] K. L. G. Tan, "Confronting Cyberterrorism with Cyber Deception," December 2003. [Online]. Available: http://www.au.af.mil/au/awc/awcgate/nps/cyberterr_cyberdecep.pdf.

[157] G. Stoneburner, C. Hayden and A. Feringa, "Engineering Principles for Information Technology Security (A Baseline for Achieving Security), NIST Special Publication 800-27 Rev A," June 2004. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-27A/SP800-27-RevA.pdf.

[158] G. G. Preckshot, "Method for Performing Diversity and Defense-in-Depth Analyses of Reactor Protection Systems," December 1994. [Online]. Available: http://pbadupws.nrc.gov/docs/ML0717/ML071790509.pdf.

[159] M. MacQueen and W. Boyer, "Deception used for Cyber Defense of Control Systems," in *Proceedings of the 2009 IEEE Conference on Human-System Interactions (HSI 2009)*, Catania, Italy, 2009.

[160] X. Li, C. Jiang, J. Li and B. Li, "VMInsight: Hardware Virtualization-based Process Security Monitoring System," in *Proceedings of the 2011 International Conference on Network Computing and Information Security*, Guilin, China , 2011.

[161] J. Kruse, S. Landsman, P. Smyton, S. Chin, A. Cooper, A. Dziewulski, H. Hawley and M. King, "POET: Integrating Political, Operational, Economic, and Technical Factors into Systems Engineering (PR Case No. 11-1825)," 2011. [Online]. Available: http://www.mitre.org/work/areas/research/2011iebriefings/20MSR058-CB.pdf.

[162] L. J. Gueck, "NRO Mission Architecture," in *Proceedings of the 2011 Ground System Architectures Workshop*, Los Angeles, CA, 2012.

[163] DoD Deputy CIO, "DoD Architectural Framework Version 2.02," August 2010. [Online]. Available: http://dodcio.defense.gov/Portals/0/Documents/DODAF/DoDAF_v2-02_web.pdf.

[164] Assistant Secretary of Defense for Research and Engineering (ASD(R&E)), "Technology Readiness Assessment (TRA) Guidance," 13 May 2011. [Online]. Available: http://www.acq.osd.mil/ddre/publications/docs/TRA2011.pdf.

Appendix A Mapping Cyber Resiliency Techniques to Sub-Objectives and Objectives

The following table provides additional detail on mapping cyber resiliency techniques to sub-objectives (and thereby to objectives), by identifying cyber defender activities enabled by the techniques, activities that help achieve sub-objectives.

Table 12. Detailed Mapping of Cyber Resiliency Techniques to Sub-Objectives and Objectives

Objective	Sub-Objective	Activity	Technique
Understand: *maintain useful representations of mission/business cyber dependencies, and of the status of cyber resources with respect to possible adversary activities*	Understand adversaries	Share threat information	Supporting
		Perform or support malware and forensic analysis	Analytic Monitoring
		Perform retrospective analysis to investigate historical trends and activities	Analytic Monitoring
		Observe and analyze adversary activities in deception environments	Deception & Analytic Monitoring
	Understand mission or business function dependencies resources *and* Understand the functional dependencies among cyber resources	Perform mission impact analysis, business impact analysis, or crown jewels analysis to identify critical, essential, and supporting assets / capabilities	Supporting
		Identify, and maintain a representation of, functional and mission dependencies among cyber resources	Dynamic Representation
		Identify mission / business function dependencies on cyber resources	Realignment
		Identify non-mission / business function dependencies on or uses of cyber resources	Realignment
		Identify dependencies and interactions among cyber defenses, security controls, and performance controls	Coordinated Defense
		Determine degrees of criticality of cyber resources, thereby identifying critical assets	Privilege Restriction
		Determine types and degrees of trust for users and cyber entities (e.g., components, data, processes, interfaces)	Privilege Restriction
		Define an implementable set of change parameters (e.g., conditions under which unpredictable changes should not be made, “distance” beyond which a service should not be moved, ranges for frequency of changes)	Unpredictability

Objective	Sub-Objective	Activity	Technique
	Understand the status of resources with respect to adversary activities	Track security posture of cyber resources (e.g., patch status, compliance with configuration guidance)	Supporting
		Track effectiveness of CCoA and adapt as necessary	Adaptive Response
		Coordinate sensor coverage to avoid gaps or blind spots	Analytic Monitoring
		Correlate or otherwise combine data from different sensors	Analytic Monitoring
		Detect by analyzing data to identify anomalies, develop I&W, and assess likelihood of compromise or intrusion	Analytic Monitoring
		Analyze data to assess lifespan / retention conditions for Non-Persistence	Analytic Monitoring
		Analyze data to assess effectiveness of CCoAs (in support of response)	Analytic Monitoring
		Dynamically reconfigure sensors	Analytic Monitoring
		Perform damage assessment [to understand status of resources, to support reconstitution]	Analytic Monitoring
		Dynamically relocate sensors	Dynamic Positioning
		Define and maintain a representation of the resiliency posture (including security posture, performance with respect to SLAs or KPPs, and quality as determined using Substantiated Integrity mechanisms) of cyber resources and adversary activities against cyber resources	Dynamic Representation
		Validate data provenance	Substantiated Integrity
		Validate data integrity / quality to ensure it has not been corrupted	Substantiated Integrity
		Validate software / service integrity / behavior to ensure it has not been corrupted	Substantiated Integrity
Prepare: *maintain a set of realistic cyber courses of action that address predicted cyber attacks*	Create and maintain cyber courses of action	Define / maintain realistic CCoAs, i.e., CCoAs that can be executed in a coordinated way given existing controls and management responsibilities	Coordinated Defense
	Maintain the resources needed to accomplish cyber courses of action	Back up data needed to restore or reconstitute mission and supporting functionality	Redundancy
		Maintain gold copies of mission-essential software and configuration data	Substantiated Integrity

Objective	Sub-Objective	Activity	Technique
		Provide mechanisms and/or procedures for snapshotting or otherwise capturing, and then restoring, state	Supporting
		Provide mechanisms and/or procedures for capturing and archiving data (or otherwise maintaining chain of evidence)	Supporting
	Validate the realism of cyber courses of action	Identify dependencies and interactions among cyber defenses, security controls, and performance controls	Coordinated Defense
		Simulate and/or exercise CCoAs	Dynamic Representation
Prevent: *preclude successful execution of an attack on a set of cyber resources*	Harden resources based on adversary capabilities	Define sets of cyber resources with clear boundaries	Segmentation
		Identify key locations to place mechanisms	Coordinated Defense
		Provide protection mechanisms at different locations	Coordinated Defense
		Coordinate ongoing management	Coordinated Defense
		Coordinate definition and assignment of privileges	Coordinated Defense
	Deflect adversary actions	Create and maintain deception environment(s), e.g., honeypots, honeynets, decoy documents or data stores	Deception
		Redirect adversary activities to deception environment(s)	Deception
	Dissuade / deter adversaries by increasing the adversary's costs	Maintain or dynamically create determinably different instantiations / implementations of capabilities or component functionality (e.g., different operating systems, applications, hardware)	Diversity
		Define and maintain determinably different alternative processing paths (i.e., different sequences of services or applications used to respond to the same request)	Diversity
		Define and maintain determinably different alternative communications paths (e.g., different protocols, different communications media)	Diversity
		Use determinably different supply chains for key technical components	Diversity
		Identify and maintain determinably different mission data sources	Diversity
		Create and maintain determinably different information stores	Diversity

Objective	Sub-Objective	Activity	Technique
		Assign privileges based on types and degrees of trust	Privilege Restriction
		Assign and maintain privilege restrictions, particularly for critical assets	Privilege Restriction
		Define enclaves or sets of cyber resources with clear boundaries	Segmentation
		Maintain boundary protections	Segmentation
		Reconfigure components and services, use alternative equivalent components or services, or dynamically reposition processing randomly, in accordance with change parameters	Unpredictability
	Dissuade / deter adversaries by increasing the adversary's risks	Reveal adversary TTPs by analysis	Analytic Monitoring
		Cause the adversary to reveal TTPs by directing adversary activities to a deception environment	Deception
	Deter attacks by limiting the adversary's benefits	Conceal mission processing and communications, e.g., function hiding	Deception
		Transform data for obfuscation	Deception
		Identify services and information to which non-persistence can be applied	Non-Persistence
		Define lifespan conditions for services and connectivity	Non-Persistence
		Terminate services or connectivity when lifespan conditions no longer hold	Non-Persistence
		Define retention conditions for information	Non-Persistence
		Delete or move information when retention conditions no longer hold	Non-Persistence
		Provide mechanisms and/or procedures for capturing and archiving data (or otherwise maintaining chain of evidence)	Supporting
		Ensure that termination, deletion, or movement does not leave residual data or software behind	Non-Persistence
Continue: *maximize the duration and viability of essential mission/business functions during an attack*	Maintain functioning	Select and tailor CCoA	Adaptive Response
		Dynamically reconfigure existing resources	Adaptive Response
		Dynamically provision by reallocating existing resources	Adaptive Response
		Dynamically reconstitute critical assets or capabilities	Adaptive Response

Objective	Sub-Objective	Activity	Technique
		Deploy diverse resources rapidly (e.g., in near real time)	Diversity
		Coordinate response activities to ensure synergy rather than interference	Coordinated Defense
		Fail over to replicated resources	Supporting
		Track effectiveness of CCoA and adapt as necessary	Adaptive Response
	Ensure that functioning is correct	Validate data provenance	Substantiated Integrity
		Validate data integrity / quality to ensure it has not been corrupted	Substantiated Integrity
		Validate software / service integrity / behavior to ensure it has not been corrupted	Substantiated Integrity
	Extend the surface that an adversary must attack to be successful	Distribute mission / business functions	Realignment
		Terminate services or connectivity when lifespan conditions no longer hold	Non-Persistence
		Extend the attack surface to privilege management and privilege enforcement mechanisms	Privilege Restriction
		Reconfigure components and services, use alternative equivalent components or services, or reposition processing randomly, in accordance with change parameters	Unpredictability & Dynamic Positioning
Constrain: *limit damage from an adversary's attacks*	Isolate resources to preclude or limit adversary access	Isolate the enclave or set of cyber resources to contain adversary activities	Segmentation
	Move resources to preclude adversary access	Use distributed processing and virtualization to relocate targeted resources	Dynamic Positioning
		Dynamically relocate critical assets	Dynamic Positioning
		Reassign / relocate non-critical assets to reduce the exposure of critical assets to compromised non-critical assets	Realignment
	Change or remove resources to limit or preclude adversary access	Terminate services or connectivity when lifespan conditions no longer hold	Non-Persistence
		Delete or move information when retention conditions no longer hold	Non-Persistence
		Ensure that termination, deletion, or movement does not leave residual data or software behind	Non-Persistence
		Increase or decrease privilege restrictions based on adversary activities	Privilege Restriction

Objective	Sub-Objective	Activity	Technique
		Dynamically reconfigure critical assets	Adaptive Response
Reconstitute: *redeploy cyber resources to provide as complete a set of mission / business functionality as possible subsequent to a successful attack*	Maintain deployable / redeployable resources	Maintain multiple protected instances of hardware	Redundancy
		Create and maintain multiple protected instances of software	Redundancy
		Create and maintain multiple protected instances of information	Redundancy
	Restore functionality	Execute recovery procedures in accordance with contingency or continuity of operations plans	Supporting
		Dynamically reconstitute critical assets or capabilities	Adaptive Response
		Identify and restore non-critical functional capabilities	Adaptive Response
		Coordinate recovery activities to avoid gaps in security coverage	Coordinated Defense
	Validate functionality	Validate data provenance	Substantiated Integrity
		Validate data integrity / quality to ensure it has not been corrupted	Substantiated Integrity
		Validate software / service integrity / behavior to ensure it has not been corrupted	Substantiated Integrity
Transform: *change aspects of organizational behavior in response to prior, current, or prospective adversary attacks*	Identify unnecessary dependencies	Assess mission / business function risks due to dependency on resources shared with non-mission functions	Realignment
	Adapt systems and mission / business processes to mitigate risks	Reallocate resources and/or reassign administrative / management responsibility based on risk to mission / business function	Realignment
		Identify and remove or replace data feeds and connections for which risks outweigh benefits	Realignment
Re-Architect: *modify architectures for improved resiliency*	Address predicted long-term changes in adversary capabilities, intent, and/or targeting	Integrate cyber resiliency strategy with other organizational strategies	Supporting
	Apply cyber resiliency practices cost-effectively	Identify and weigh alternatives for applying cyber resiliency practices as part of systems engineering	Supporting
	Incorporate emerging technologies in ways that improve (or at least do not degrade) cyber resiliency	Analyze emerging technologies with respect to cyber resiliency	Supporting

Appendix B Initial Questionnaire

A cyber resiliency assessment uses value scales for objectives, sub-objectives, and techniques to provide qualitative assessments and develop recommendations. Those scales assume a level of understanding of the cyber resiliency engineering framework beyond the general understanding that can be fostered via an introductory briefing. Thus, a cyber resiliency assessment involves interviews and document review by resiliency experts to obtain the information necessary to make an assessment, rather than using an instrument that would ask multiple stakeholders to make assessments. A structured questionnaire can support information gathering from Program Managers, architects and systems engineers, system managers and administrators, and cyber defenders. The questions follow the flow of the cyber resiliency objectives. For each objective, general questions may be included, followed by questions related to the use of cyber resiliency techniques. Note that the interviewer(s) will tailor the questionnaire, both prior to the interview, based on the purpose and scope of the cyber resiliency assessment and on review of relevant documentation, and during the interview based on responses. Note also that different interviewees will be able to answer different questions.

Questions
Objective: Understand
General
G.1 How important is it to understand your adversary? Do you use intelligence (which could consist of information from commercial security service providers) to inform your planning and response?
G.2 Do you have a view of the nature of your adversary? Who is the typical adversary you guard against? 1. Do you have a view or estimate of the adversary's capabilities and/or tradecraft? 2. Do you have a view of the adversary's intent? 3. Do you have a view of what the adversary may be targeting and how persistent they will be in their efforts?
G.3 Have you captured your view of the above in any written/formalized way (e.g., in a risk assessment report)?
G.4 Have you completely identified all of your cyber assets, resources, and capabilities? How current is your inventory? What mechanisms and/or processes do you use to maintain your inventory?
G.5 What is your understanding of your mission or your key business function dependencies? 1. Do you have a mapping, listing, or other representation of which cyber resources are essential to, critical to, or supportive of specific missions/business functions? 2. Do you have a mapping, listing, or other representation of dependencies among cyber resources?
G.6 Can you identify dependencies and interactions among cyber defenses, security controls, and performance controls?
G.7 Have you captured this understanding in any written/formalized way (e.g., a contingency or continuity of operations plan)?
G.8 Do you have any way of checking the status of your cyber resources with respect to your adversary (e.g., how vulnerable they are, whether they have been compromised, etc.)?
Technique: Analytic Monitoring
Are you able to determine whether a cyber attack is underway?
Are you able to observe/monitor adversary activities? If so, how do you do this (e.g., what sensors do you use, where are they placed, what can they collect, what do they ordinarily collect)?
Do you retain the data you capture for forensic analysis?
Do you perform or support malware and forensic analysis? If not, do you employ or send your information to some other party that does perform such analysis?
Do you perform retrospective analysis to investigate historical trends and activities?
Do you have internal sensors to monitor intrusions by the adversary?

Questions
If you have identified mission / business function dependencies on cyber resources (see question G.5), can you calibrate sensors based on resource criticality?
Have you defined conditions under which unpredictable changes should not be made, “distance” beyond which a service should not be moved, ranges for frequency of changes? If so, do you have a way to monitor for changes that exceed the bounds you have set?
Do you track the security posture of cyber resources (e.g., patch status, compliance with configuration guidance)?
Do you track the effectiveness of cyber courses of actions and adapt as necessary?
Technique: Deception
One way to learn more about an adversary and their tradecraft is to employ deception environments (e g., honeynets). Do you take any action to observe the adversary in a deception environment?
Have you or do you have plans to perform damage assessments [to understand status of resources, to support reconstitution]?
Technique: Dynamic Positioning
Do you dynamically reposition or reconfigure sensors?
Technique: Substantiated Integrity
In general, do you monitor the integrity of data received and stored (e.g., spot-checks, in/outflow checks, etc.)?
Are you able to determine and validate the origin of the data you receive (so that you know it is genuine)?
Do you validate data integrity / quality to ensure it has not been corrupted? How do you do so?
Do you validate software / service integrity / behavior to ensure it has not been corrupted? If so, how?
Objective: Prepare
Technique: Coordinated Defense
One means to protect your infrastructure is to employ a layered defense strategy (protection in depth). Do you use such a strategy? If so, what is involved? For example, do you use multiple IDSs or AVs?
What actions (if any) do you take to “harden” your cyber resources?
Have you defined and do you maintain cyber courses of action to take in response to an attack? Have you tested them in a coordinated way?
How often do you back up data needed to restore or reconstitute mission and supporting functionality?
Do you place your protection mechanisms at different locations within your infrastructure?
How do you identify key locations to place mechanisms?
Do you take actions to coordinate your defensive measures so they do not interfere with each other?
Do you perform any simulations, exercises or war-gaming of how the cyber course of action you currently employ would respond to adversary actions? If so, how do you employ the results of such activities?
Technique: Privilege Restriction
Can you determine types and degrees of trust for users and cyber entities (e.g., components, data, processes, interfaces)?
Do you assign privileges? Do you use a method that depends on types and degrees of trust?
Do you assign and maintain privilege restrictions, particularly for critical assets?
Technique: Deception
Do you create and have deception environments (e.g., honeypots), or place decoy documentation on accessible data stores?
If you maintain deception environments, what actions, if any, do you take to deflect adversaries to those environments?
Objective: Prevent
General
G.9 What actions do you take to dissuade your adversaries from attacking?
G.10 What actions do you take to make the adversaries work harder at achieving success?

Questions
Technique: Diversity
Do you maintain different implementations of capabilities or component functionality (e.g., different operating systems, applications, hardware)?
Do you maintain different alternative processing paths (i.e., different sequences of services or applications used to respond to the same request)?
Do you maintain determinably alternative communications paths (e.g., different protocols, different communications media)?
Does your critical data come from a single source or do you employ multiple data sources for your critical data?
Technique: Segmentation
Have you defined enclaves or sets of cyber resources with a clear boundary within your intranet?
Do you maintain boundary protections within your intranet? If so, how (e.g., logical isolation, physical isolation, a combination)?
Do you isolate your security operations center from the rest of your network?
Technique: Unpredictability
Do you take any actions to make it more difficult for an adversary to predict the behavior of your systems, networks, or users? If so please elaborate.
For example, do you reconfigure components or services at some random interval?
For example, do you switch/swap services, replacing them with equivalent services, at some random interval (e.g., change browsers)?
Technique: Deception
Do you conceal key/important processing and communications, e.g., function hiding?
Do you hide/conceal your processing, functions or data? What methods do you use?
Technique: Non-Persistence
A good means of preventing the adversary of achieving a permanent foothold is to refresh data storage and processing periodically (this is known as non-persistence). Do you take any such actions?
Do you use virtualization (or some other technique) to allow you to periodically refresh services? If so, do you do so for security reasons (prevent adversary from gaining a foothold) or some other reason?
If you do employ non-persistence, do you set some lifespan conditions (thresholds) for refreshing or terminating services?
Do you terminate services or connections when they are no longer needed?
Do you delete data when it is no longer needed?
Do you dynamically reconstitute critical assets or capabilities?
Objective: Continue
Technique: Substantiated Integrity
Adversaries sometimes corrupt key data or services, not to the extent of blatantly crippling services so that they fail, but in such a way that the longer the corrupted services run the greater the damage that occurs. Do you take any actions to address this threat?
What actions, if any, do you take to ensure that your operations are functioning correctly (i.e., adversary has not corrupted operations)?
What actions, if any, do you take to validate that software / service integrity / behavior has not been corrupted?
What actions, if any, do you take to validate that data integrity has not been corrupted?
Technique: Adaptive Response
What actions, if any, do you take to continue operations even if an attack is underway? What is the relationship between cyber defense and execution of contingency plans or continuity of operations plans?
Do you dynamically reconfigure resources in any way so that the configuration is different than what the adversary anticipated?
Do you assess the effectiveness of your countermeasures during an attack?

Questions
Technique: Diversity
If all the systems in your organization use the same OS and applications in the same configurations, and if they have some underlying flaws (e.g., unknown flaw) then a successful zero day attack can take them *all* out. One means to address this is to use a variety of systems or services. While this will not stop all the attacks, it will mean that to compromise all the systems the adversary will have to use different attack techniques. Do you use diversity or heterogeneity as a defensive strategy? If so, for which components (e.g., OSs, hardware, communications paths or media) do you maintain different instances?
Your organization probably experiences some degree of *incidental* diversity, e.g., as a result of components procured at different times, due to a BYOD strategy. What forms of incidental diversity are present? Do your attack detection or response plans take advantage of incidental diversity? If so, how?
Do you deploy different configurations of systems or services as the attack is ongoing to confuse or impede the adversary?
Objective: Constrain
Technique: Segmentation
Do you take any actions resources to preclude or limit access by an adversary that has already compromised other system resources?
Do you isolate enclaves/subnets from other portions of your network? If so, how?
Do you have the ability to dynamically isolate some portions of your network from others in the event of an attack? If so, how, and do cyber courses of action include rules or guidance on when to do so?
Technique: Dynamic Positioning
If you can relocate your key services, it makes it harder for the adversary to successfully compromise them (or at the least it increases the chance that the adversary will be discovered trying to attack them). Do you take any action to move/relocate likely targets of the adversary?
If so, do you use virtualization and/or distributed processing to support such relocation?
Do you dynamically relocate critical information?
Technique: Non-Persistence
Do you take actions to remove critical information that is no longer needed so it is not accessible to adversary during an attack?
If so, how do you ensure the information is no longer accessible?
Objective: Reconstitute
Technique: Redundancy
Do you maintain multiple instances of your critical hardware? If so, are they protected to the same extent as the primary instance?
Do you maintain multiple instances of your critical software? If so, are they protected to the same extent as the primary instance?
Do you maintain multiple instances of your critical information? If so, are they protected to the same extent as the primary instance?
Technique: Substantiated Integrity
Do you provide mechanisms and/or procedures for snapshotting or otherwise capturing, and then restoring, state?
What actions do you take to ensure that restored data has not been corrupted?
What actions do you take to ensure that restored services or software have not been corrupted?
Objective: Transform
Technique: Realignment
Do you change any aspects or your organizational behavior to make your organization a less identifiable or appealing target?
Have you identified any data feeds/connections whose risks may outweigh their benefits? If so, what have you done about them?

Questions
Have you looked to see whether your organization's mission critical functions are overly dependent upon some non-mission critical functions or on functions that you do not control?
Have you done any analysis to identify any non-missions critical resources which you employ that may prove to be an entrée for an adversary?
Objective: Re-Architect
General
Have you made or are you planning to make changes to your architecture to address predicted long-term changes in an adversary's capability?
Have you made, or are you planning to make, changes to your architecture that will allow you to respond more dynamically to adversary attacks?
Technique: Diversity
Have you considered, or do you currently employ, different suites of operating systems, applications, browsers, networking protocols to make it more difficult for an adversary to successfully compromise your architecture? If so, please elaborate.
Have you considered using virtualization to allow you to more dynamically change OSs, applications, or locations of servers or services?
Technique: Non-Persistence
Have you considered, or do you currently employ, refreshing of your services so as to limit an adversary's ability to successfully establish and maintain a foothold in your system?
Technique: Dynamic Positioning
Have you considered, or do you currently employ, any means to relatively dynamically re-locate critical data repositories or services to other locations?

Appendix C Assessment Scales

This appendix provides value scales for all cyber resiliency objectives and sub-objectives, as well as for a representative set of cyber resiliency techniques and activities.

C.1 Assessment Scales for Cyber Resiliency Objectives and Sub-Objectives

Table 13 defines value scales for assessing or scoring how well cyber resiliency objectives are met, while Table 14 defines value scales for sub-objectives.

Table 13. Assessment Scales for Cyber Resiliency Objectives

Objective	Key Differentiators	Value Scales
Understand: *maintain useful representations of mission/business cyber dependencies, and of the status of cyber resources with respect to possible adversary activities*	• Depth of understanding, for o Mission and functional dependencies on cyber resources o Functional interdependencies among resources o Adversary capabilities, intent, and targeting, as evident or implicit in TTPs • Breadth of understanding, for o Dependencies across missions and business functions o Relationships (e.g., synergies, dependencies) among adversary activities • Currency of understanding, for o Mission and functional dependencies on cyber resources o Adversary activities and evolving capabilities, intent, and targeting o Security posture of cyber resources with respect to adversary activities • Types of decisions and/or planning that understanding is sufficient to understand	**High**: A deep, broad, current, and forward-looking understanding is able to inform both operational decisions (mission courses of action, cyber courses of action (CCoAs), resource allocation, supporting business processes such as maintenance and updates) and planning (for missions and mission evolution, as well as for cyber defense, architectural evolution, and technology investments). **Medium**: A relatively deep, broad, and current understanding is able to inform development, selection, and execution of CCoAs, as well as some investment planning. **Low**: A shallow and relatively static understanding, relying primarily on analytic models and tools that do not consider the advanced threat, is able to inform contingency planning and execution.
Prepare: *maintain a set of realistic cyber courses of action that address predicted cyber attacks*	• Depth of preparation, for contingencies that involve multiple events and thus require coordinated responses as well as CCoAs with conditional alternatives • Breadth of preparation, for o Ranges of contingencies o Contingencies that affect multiple missions or organizations • Realism of preparation, reflecting	**High**: A set of sophisticated (deep, broad, well-resourced) CCoAs has been established, kept current, and extensively validated. CCoAs are coordinated with mission CoAs. **Medium**: A set of CCoAs, including some which are broad and/or deep, has been established and validated. Resources are maintained to support CCoA execution.

Objective	Key Differentiators	Value Scales
	o Current and anticipated mission and functional dependencies on cyber resources o Maintenance of cyber resources and alternatives that are actually available o Exercises that enable CCoAs to be tried out and validated	**Low**: Cyber courses of action have been established for a few common situations (e.g., malware infection, compromise of a user or an administrator account, security spills). All other contingencies are treated as under the purview of COOP.
Prevent: *preclude successful execution of an attack on a set of cyber resources*	• Depth of prevention, including o Selective application of resource hardening to provide defense-in-depth o Coordination of defensive mechanisms • Breadth of approaches for dissuading, deterring, or deflecting adversary actions • Effectiveness of approaches (how effective are the approaches / mechanisms at reducing the likelihood of an attack being successful, and/or at increasing the adversary's work factor?)	**High**: A broad range of mechanisms, applied strategically and in depth, significantly reduce the likelihood that adversary activities will affect mission tasks that depend on cyber resources, and/or significantly increase the adversary's work factor.
		Medium: A variety of mechanisms, applied with varying degrees of depth, reduce the likelihood that adversary activities will affect mission-critical functions that depend on cyber resources, and/or increase the adversary's work factor.
		Low: Multiple mechanisms are applied with the goal of reducing the adversary's likelihood of success and/or increasing the adversary's work factor.
Continue: *maximize the duration and viability of essential mission / business functions during an attack*	• Scope of continuity o Criticality of functions (mission-critical, mission-essential) o Whether the current mission is taken into consideration • Viability (correctness, reliability) of functions • Extent to which adversary work factor is increased by changes in the attack surface	**High**: Mission-essential functions can be performed, with degradation and periods of interruption well within specified requirements, and with a high degree of confidence in the correctness or reliability of cyber resources, until the current mission is completed. Dynamic changes in the attack surface significantly increase the adversary's work factor.
		Medium: Mission-critical functions can be performed, with degradation and periods of interruption within specified requirements, and with a significant degree of confidence in the correctness or reliability of cyber resources, until the current mission is completed. Dynamic changes in the attack surface increase the adversary's work factor.
		Low: Mission-critical functions can be supported in accordance with contingency or continuity of operations planning.

Objective	Key Differentiators	Value Scales
Constrain: *limit damage from an adversary's attacks*	• Type of constraint / damage limitation o Physical / logical – minimize set of resources exposed to adversary attacks o Temporal – minimize period of time for which resources are exposed to adversary attacks • Degree of forethought / planning / support (ad-hoc, established, well-supported by automation) • Types of resources to which exposure to adversary attacks can be limited o Physical (e.g., subnets, platforms) o Logical (e.g., processes, databases or data stores, connections)	**High**: Damage from adversary attacks can be limited quickly, temporally as well as physically / logically, and to relatively few resources beyond those affected when the attack was identified or indicated, based on established and well-supported procedures. Damage limitation applies to a wide range of physical and logical resources.
		Medium: Damage from adversary attacks can be constrained temporally and/or physically / logically, via manual intervention to isolate resources and/or to restart some processes, using established procedures.
		Low: Damage from adversary attacks can be somewhat constrained, primarily via manual intervention to isolate physical resources.
Reconstitute: *redeploy cyber resources to provide as complete a set of mission/business functionality as possible subsequent to a successful attack*	• Effectiveness of reconstitution / extent of mission support • Speed of reconstitution • Confidence in reconstitution o Correctness / compliance with functional or quality requirements o Completeness ▪ Set of restored resources (e.g., services, software, hardware, databases or other structured data stores) ▪ Period of time for which mission data is missing (unknown, known, minimal) o Integrity / validation that reconstituted resources have not been corrupted	**High**: Functionality is reconstituted rapidly (consistent with overall operations tempo for supported missions and cyber defense operations) by redeploying cyber resources to provide a full range of mission support, and to a known level of confidence in correctness, completeness, and integrity.
		Medium: Functionality is reconstituted by redeploying cyber resources to provide mission-essential and mission-critical functions, with some confidence that they have not been corrupted.
		Low: Functionality is reconstituted by redeploying cyber resources in accordance with contingency planning or COOP.
Transform: *change aspects of organizational behavior in response to prior, current, or prospective adversary attack*	• Integration of the strategy of reducing mission dependencies on non-mission-critical assets into organizational processes o Contingency planning o Mission/business process re-engineering o Systems engineering & architectural development • Capabilities or functions integrated into organizational processes	**High**: Attack surface reduction is a key element of the organization's systems engineering / architectural development, and is integrated with process re-engineering and contingency planning. The organization reduces/eliminates as many dependencies on non-mission critical services/assets as feasible, and eliminates, to the extent feasible, non-mission essential behavior.

Objective	Key Differentiators	Value Scales
	o Identification/specification of organizational behavior ▪ Of different types (normal behaviors, mission-essential behaviors, mission-critical behaviors) ▪ Under different contingencies (normal, stressed/hazard, cyber attack) o Identification of exposed services / resources	**Medium**: The organization considers attack surface reduction as part of systems engineering, while reduction of non-essential dependencies is considered as part of process re-engineering and contingency planning. The organization reduces / eliminates some dependencies on non-mission critical services / resources, and eliminates some non-mission essential behavior.
		Low: As part of contingency planning, the organization identifies critical and non-critical services / resources and mission essential organizational-behavior. For critical resources, organization identifies those dependencies that such resources have with non-critical resources.
Re-Architect: *modify architectures to improve resiliency*	• Integration of cyber resiliency principles with architectural roadmap development o Consideration of possible adversary TTPs in architectural definition o Support for dynamic mechanisms (e.g., using virtualization or other technologies that enable rapid changes in the attack surface) o Support for emerging cyber resiliency technologies and solutions • Investment in custom components via re-implementation or custom development o Software o Hardware	**High**: Architectural roadmap development considers not only possible adversary TTPs, but the fact that adversaries will develop unforeseen TTPs and campaign strategies. The architectural roadmap includes support for dynamic mechanisms and for cyber resiliency technologies and solutions emerging from basic and applied research. The architectural roadmap includes re-implementation of critical software and/or custom development of key critical hardware components.
		Medium: Architectural roadmap development considers possible adversary TTPs, and to some extent the fact that adversaries will develop unforeseen TTPs and campaign strategies. The architectural roadmap includes limited support for dynamic mechanisms and for cyber resiliency technologies and solutions emerging from applied research. The architectural roadmap considers re-implementation of critical software and/or custom development of key critical hardware components.
		Low: Architectural definition and planning take into consideration possible adversary TTPs.

Table 14. Assessment Scales for Cyber Resiliency Sub-Objectives

Sub-Objective	Low	Medium	High
Understand adversaries [Understand]	Adversaries are characterized in terms of general types or goals (e.g., hackers, organized crime, industrial espionage, nation-state espionage, terrorists, military adversaries). Adversary TTPs are characterized in terms of general attack methods. Adversary characterization is typically by reference to published threat reports.	Adversaries are identified in general terms and characterized in terms of capabilities, intent, and targeting. Adversary TTPs are identified in terms of specific attacks or exploitation techniques (for example, using CAPEC, http://capec.mitre.org/). Information about adversary activities is refreshed on a regular basis, based on analysis of monitoring data and/or on intelligence information supplied by external entities.	Adversaries are identified specifically, based on analysis of observed attacks. Adversaries are described in terms of capabilities, specific goals or objectives, and targeted cyber resources, providing the foundation for an adversary model. Adversary TTPs are identified in terms of specific attacks or exploitation techniques. Information about adversary activities is refreshed on an ongoing basis, based on analysis of monitoring data and/or on intelligence information supplied by external entities. An adversary model enables assessment of the adversary work factor associated with a cyber course of action or an architectural alternative.
Understand mission or business function dependencies [Understand]	Dependencies of mission or business functions on cyber resources are represented in a static, manually-generated representation, typically as part of design or COOP documentation.	Dependencies of mission or business functions on cyber resources are identified (to system / infrastructure providers, to contingency planners, and to a lesser extent to cyber defenders) via a combination of static and dynamic representations, supported by automation (e.g., Business Service Management (BSM) or IT Service Management (ITSM) tools).	Cyber defenders are provided with a dynamic representation of mission / business function dependencies on cyber resources, automatically updated based on observation and analysis of usage. System / infrastructure providers and others responsible for COOP execution are presented with snapshots as needed to support planning.
Understand status of resources with respect to adversary activities [Understand]	Security posture monitored at different architectural layers, e.g., using basic intrusion detection and malware protection.	Security posture monitored at the enterprise level, using continuous monitoring and risk scoring tools and processes, which integrate the results of security posture monitoring at different architectural layers.	Security posture monitored at the enterprise level, using continuous monitoring and risk scoring tools and processes, which integrate the results of security posture monitoring at different architectural layers. Capabilities are provided and used that enable the security posture of specific resources to be monitored more closely, in response to I&W.

Sub-Objective	Low	Medium	High
Create and maintain cyber courses of action [Prepare]	Cyber courses of action have been established for a few common situations (e.g., malware infection, compromise of a user or an administrator account, security spills).	Cyber courses of action have been established based on an understanding of anticipated adversary activities. CCoAs take into consideration mission dependencies on resources, and the effects of defender actions on mission capabilities; thus, CCoAs are updated as overall mission needs change as well as in response to changing threat information.	Cyber courses of action have been established based on an understanding of anticipated and potential adversary activities, as well as on anticipated changes in systems and dependencies. CCoAs take into consideration mission dependencies on resources, and the effects of defender actions on mission capabilities; thus, CCoAs are updated dynamically as specific mission needs change as well as in response to changing threat information. CCoAs take into consideration the adversary model, in particular the adversary work factor associated with different CCoAs.
Maintain resources needed to accomplish cyber courses of action [Prepare]	Gold copies of mission-essential software, configuration data, and (as relevant) static mission data are retained. Mission data is backed up as required by COOP.	Gold copies of mission-essential software and configuration data are kept current with patches and configuration changes. Mission data is backed up or checkpointed consistent with mission requirements. Redundant resources (e.g., hardware, alternative software, storage, communications) are maintained, consistent with diversity and redundancy approaches.	Gold copies of mission-essential software and configuration data are kept current with patches and configuration changes. Mission data is backed up or checkpointed consistent with mission requirements. Mechanisms are provided for snapshotting, checkpointing, or capturing, and restoring, state. Redundant resources (e.g., hardware, alternative software, storage, communications) are maintained, consistent with diversity, redundancy, and dynamic positioning strategies.
Validate the realism of cyber courses of action [Prepare]	Gold copies are periodically validated (e.g., the patch status of gold copies of software is checked, backup-and-restore processes are exercised), typically as part of COOP exercises.	Cyber courses of action are exercised as part of cyber defense exercises, and are integrated into COOP exercises. CCoA exercises include validation of key software and data.	CCoAs are exercised as part of mission as well as cyber defense exercises. Mission impacts of CCoAs are evaluated, in light of observed mission dependencies (indirect as well as direct) on cyber resources.

Sub-Objective	Low	Medium	High
Harden resources based on adversary capabilities[19] as well as resource criticality [Prevent]	Resources hardening (e.g., security controls tailored or supplemented above baselines, configurations controlled more tightly) is based on general assumptions about adversary capabilities.	Resource hardening is based on an assessment of adversary capabilities to discover, discover vulnerabilities in, and exploit vulnerabilities in the resources. Defensive mechanisms are coordinated to avoid undesirable interactions.	Resource hardening is based on threat intelligence about adversary capabilities to discover, discover and/or implant vulnerabilities in, and exploit vulnerabilities in the resources. Defensive mechanisms are coordinated to avoid undesirable interactions, provide synergy as feasible, and facilitate monitoring and analysis.
Deflect adversary actions[20] [Prevent]	A deception environment (e.g., a honeypot on a DMZ) is established.	A deception environment is established and maintained (e.g., a honeynet with information updated on a periodic basis so that information appears fresh); capabilities are provided for cyber defenders to deflect suspect network traffic to the deception environment.	A sophisticated deception environment is established and maintained (e.g., a honeynet with information and configurations updated on a periodic basis and in response to observable mission activities so that information appears fresh and realistic); capabilities are provided for cyber defenders to deflect suspect network traffic to the deception environment and monitor adversary behavior.
Dissuade / deter adversaries by increasing the adversary's costs [Prevent]	Mechanisms from at least one cyber resiliency technique intended to increase the adversary's costs (diversity, segmentation, privilege restriction, unpredictability) are employed; increases in the adversary's costs are assumed to track the number of specific mechanisms.	Mechanisms from one or more cyber resiliency techniques intended to increase the adversary's costs are employed; increases in the adversary's costs are assumed to track the number of specific mechanisms.	A multi-faceted resiliency strategy employs multiple techniques for increasing the adversary's costs. The strategy is informed by an adversary model that enables the adversary's costs for alternative CCoAs or architectural alternatives to be assessed, and that takes into consideration the adversary's concern for revealing TTPs as well as the adversary's specific targets or goals..

[19] The assessment of this sub-objective should be less than or equal to the assessment of "Understand adversaries."
[20] This sub-objective is relevant only if deception is used.

Sub-Objective	Low	Medium	High
Dissuade / deter adversaries by increasing the adversary's risks[21] [Prevent]	Analytic Monitoring mechanisms are used to increase the likelihood that adversary TTPs will be revealed.	Analytic Monitoring and some limited Deception mechanisms are used to obtain information about adversary TTPs.	A multi-faceted resiliency strategy employs multiple techniques for increasing the adversary's risks. The strategy is informed by an adversary model that takes into consideration the adversary's concern for revealing TTPs, as well as for attribution.
Deter attacks by limiting the adversary's perceived benefits[22] [Prevent]	Data targets are hidden or obfuscated.	Mechanisms from one or more cyber resiliency techniques intended to limit adversary's perceived benefits (e.g., deception, non-persistence) are employed.	A multi-faceted resiliency strategy employs multiple techniques for limiting the adversary's perceived benefits.
Maintain functioning [Continue]	Contingency or continuity of operations plans provide procedures for graceful degradation and/or failing over to replicated resources. Such cyber courses of action as have been established focus on immediate remediation or clean-up (e.g., for security spills or malware infection) rather than on mission continuity.	CCoAs enable mission-critical functions to be performed until the current mission is completed, with degradation and periods of service interruption within specified requirements. Resources to support CCoAs enable dynamic reconfiguration, provisioning, and/or reconstitution, but such actions may entail significant operator / defender intervention. CCoAs can be tailored to address immediate mission impacts and/or observed adversary actions, and adapted to reflect observed effectiveness.	CCoAs enable mission-essential functions to be performed until the current mission is completed, with degradation / periods of service interruption well within specified requirements. Resources to support CCoAs enable dynamic reconfiguration, provisioning, and/or reconstitution, with some degree of operator / defender intervention. CCoAs can be tailored to address observed, predicted, or inferred mission impacts; immediate mission needs and future mission implications; and observed or indicated adversary activities. CCoAs can be adapted to reflect observed effectiveness.
Ensure functioning is correct [Continue]	Correctness checking relies primarily on the use of gold copies, and on ad-hoc (and usually end-user) observation of the behavior of applications.	Cyber courses of action include use of mechanisms to check that the functioning of mission applications, systems, and networks falls within acceptable behavioral parameters.	CCoAs include use of mechanisms to validate the correctness of cyber resources, including data and software. Functional behavior is evaluated using not only acceptable behavioral parameters, but also threat information and the mission / operational environment.

[21] Revealing adversary TTPs increases either costs or risks, by limiting the future value of those TTPs to the adversary or by increasing the possibility of and confidence in attribution.
[22] Adversary benefits can be expressed, for example, in terms of the amount of sensitive data the adversary could exfiltrate.

Sub-Objective	Low	Medium	High
Extend the area an adversary must attack to be successful [Continue]	Contingency or continuity of operations plans take advantage of distribution of mission / business functions (typically using process virtualization). Changes in the attack surface are primarily the result of performance optimization (e.g., movement of virtual machines).	Cyber courses of action include use of mechanisms to change the attack surface using non-persistence and unpredictable changes, with operator / defender intervention. CCoAs take advantage of distribution of mission / business functions and/or supporting cyber resources, typically using process and storage virtualization.	CCoAs include the use of non-persistence and unpredictability in a mission-sensitive (i.e., non-disruptive) way. In addition, mechanisms for changing the attack surface unpredictably and for terminating unneeded services or connections are implemented in a mission-sensitive way. CCoAs take advantage of distribution of mission / business functions and/or supporting cyber resources, typically using process and storage virtualization as well as redundant control and information flows.
Isolate resources to preclude or limit adversary access [Constrain]	Methods for isolating resources are ad-hoc and tied to management or administrative spans of control, and primarily take the form of terminating services or connections between an enclave and a wide-area network (e.g., the Internet).	Methods for isolating resources are established and included in CCoAs, but primarily take the form of terminating services or connections between enclaves separated on the basis of information policies or technology classes (e.g., legacy vs. emerging).	Methods for isolating or limiting access to resources are established, included in CCoAs, and supported by automation. Methods include terminating services or connections to, or escalating privileges required to use, resources (including enclaves, services or applications, and knowledge bases).
Move resources to preclude adversary access [Constrain]	Methods for moving resources for which adversary activity is evident are ad-hoc and primarily take the form of terminating functioning in one location (physical or virtual) and restarting it in another, typically at the cost of some period of outage, loss of intermediate processing results, or source data loss.	Methods for moving resources for which adversary activity is evident or indicated are established and included in cyber courses of action, and include capturing state for some processes, as well as termination/restart for others.	Methods for moving resources for which adversary activity is evident or indicated are established, included in CCoAs, and supported by automation. Methods include capturing state for some processes, as well as termination/restart for others. In addition, CCoAs include moving critical resources proactively.

Sub-Objective	Low	Medium	High
Change or remove resources to limit or preclude adversary access [Constrain]	Standard operations terminate services and connectivity, and remove data from rapid-access (vs. long-term) storage, based on performance considerations.	Standard operations terminate services and connectivity, and remove data from rapid-access (vs. long-term) storage, based on predetermined lifespan considerations (which include usage and performance factors). CCoAs include established methods for removing resources and for reconfiguring critical resources.	Standard operations terminate services and connectivity, and remove data from rapid-access (vs. long-term) storage, based on predetermined lifespan considerations, which include mission priorities as well as usage and performance factors. CCoAs include established methods for removing resources, reconfiguring critical resources, and escalating privileges required to use resources.
Maintain deployable / redeployable resources [Reconstitute]	Redundant cyber resources (e.g., hardware, communications) can be deployed, and selected cyber resources (e.g., VMs) can be redeployed, as part of COOP execution, based on determination of mission criticality.	Duplicative[23] cyber resources can be deployed, and selected cyber resources can be redeployed, to support defined CCoAs. Duplicative cyber resources are maintained in an up-to-date security posture (e.g., with patches and configuration changes, with current assignments of privileges or privilege constraints).	Duplicative cyber resources can be deployed, and selected cyber resources can be redeployed, to support defined CCoAs. Duplicative cyber resources are maintained in an up-to-date security posture (e.g., with patches and configuration changes, with current assignments of privileges or privilege constraints). The determination of which cyber resources to make redundant / duplicative and/or redeployable is driven not only by mission criticality, but also by an assessment of inherent vulnerabilities.
Restore functionality [Reconstitute]	As part of COOP execution, mission-essential functions can be reconstituted in accordance with mission / functional requirements.	CCoAs enable mission-essential functions to be so reconstituted as to decrease the likelihood that a repeated or follow-on attack would affect them in the same way.	CCoAs enable mission-essential, mission-critical, and most other functionality to be so reconstituted as to decrease the likelihood that a repeated or follow-on attack would affect them in the same way.
Validate functionality [Reconstitute]	As part of COOP execution, mission-essential functioning is determined to be consistent with its prior behavior.	CCoAs for reconstitution / recovery include limited validation (e.g., via checksums, via human inspection of data) of cyber resources to ensure that they have not been corrupted.	CCoAs for reconstitution / recovery include automated validation of cyber resources to ensure that they have not been corrupted, and include actions to take in case corruption is found.

[23] Duplicative resources provide the same functionality or capabilities, but may be in a different form.

Sub-Objective	Low	Medium	High
Identify unnecessary dependencies[24] [Transform]	As part of long-term investment planning, non-essential functions are identified that could be off-loaded or outsourced. Cyber resources that are needed only to support such functions are identified; planning considers whether such resources could be repurposed.	In addition, based on an understanding of mission or business function dependencies, cyber resources are identified that support both essential and non-essential functions. Planning considers whether resources could be reallocated so that non-essential functions could be separated from essential ones.	In addition, based on an understanding of mission dependencies and of functional interdependencies, cyber resources are identified that could be high-value stepping-stone targets for adversary attacks.
Adapt systems and mission / business processes to mitigate risks[25] [Transform]	Systems and/or mission / business processes are enhanced (e.g., by additional controls or procedures) to mitigate risks due to observed or expected adversary activities, as part of an overall security risk management process.	Systems and mission / business processed are changed to mitigate risks to observed or expected adversary activities. In particular, systems are changed to reallocate cyber resources, so that non-essential functions are separated from essential ones, and to improve cyber situation awareness; mission / business processes are changed to restrict privileges.	Systems and mission / business processed are changed to mitigate risks to anticipated or potential adversary activities. In particular, systems are changed to minimize the number of high-value targets, to increase the adversary's work factor, and to improve cyber situation awareness; mission / business processes are changed to be more tightly aligned with cyber defender processes.

[24] The assessment of this sub-objective should be less than or equal to the assessment of "Understand mission or business function dependencies."
[25] The assessment of this sub-objective should be less than or equal to the assessments of "Understand adversaries" and "Understand mission or business function dependencies."

Sub-Objective	Low	Medium	High
Address predicted long-term changes in adversary capabilities, intent, and/or targeting[26] [Re-Architect]	Cyber security planning (including planning for providing resilience-related security controls) is informed by predictions about the adversary. Cyber security planning is coordinated with business continuity; cyber security is part of larger-scale risk management (e.g., coordinated management of information, IT, compliance, and business risks).	Cyber security planning (including planning for providing resilience-related security controls) and mission planning (including planning for alternative mission courses of action in case of cyber contingencies) are informed by predictions about adversary capabilities, intent, and/or targeting. The cyber security, architectural, and acquisition strategies are aligned; cyber security (including cyber resiliency) is part of enterprise risk management.	Cyber security planning (including planning for providing resilience-related security controls) and mission planning (including planning for alternative mission courses of action in case of cyber contingencies) are informed by predictions about adversary capabilities, intent, and/or targeting. The cyber security, architectural, and acquisition strategies are coordinated; cyber resilience is a central part of mission assurance strategy, which is part of the organization's mission and enterprise risk management strategies. or Full integration of cyber security and resiliency into the organization's mission assurance strategy, which is a significant part of the organization's mission and enterprise risk management strategies.
Apply cyber resiliency practices cost-effectively [Re-Architect]	The initial and support costs for applying a cyber resiliency practice to an architecture are assessed.	In addition, the effects of applying a cyber resiliency practice to an architecture are assessed in terms of changes to the cyber resiliency posture.	In addition, consequential costs and benefits to all stakeholders are taken into consideration.
Incorporate emerging technologies in ways that improve (or at least do not degrade) cyber resiliency [Re-Architect]	When emerging technologies are considered for incorporation in an architecture, potential new or additional security risks (e.g., due to additional attack paths) are identified and (as appropriate) mitigated.	In addition, challenges to cyber situational awareness (e.g., due to lack of monitoring capabilities for the technologies) are identified and addressed. CCoAs are defined and supporting technologies in the architecture are configured to enable the emerging technologies to be isolated.	In addition, CCoAs and supporting technologies in the architecture provide for validation of the emerging technologies during recovery; procedures are applied prior to and during integration of the emerging technologies into the architecture to validate its correct functioning.

C.2 Assessment Scales for Selected Cyber Resiliency Techniques

Table 8 provides a general scale for assessing how well an architecture incorporates – or supports the future incorporation of – cyber resiliency techniques. However, particularly when

[26] The descriptions of Low, Medium, and High are adapted from Table 2 of [7], with Low corresponding to Cyber Prep Level 2, Medium to Level 3, and High to Levels 4 and 5.

the assessment supports an AoA or development of specific recommendations to Program Managers or system owners, a more nuanced assessment may be needed to identify capabilities and gaps. For each technique, a set of differentiating factors has been defined in Table 7. Assessment scales for the differentiating factors can then be defined, together with representative recommendations for more effective incorporation of the technique into the architecture. Differentiating factors and definitions of assessment levels are presented below for the following techniques: Adaptive Response, Analytic Monitoring, Coordinated Response, Diversity, Privilege Restriction, and Redundancy. Definitions of assessment levels (and corresponding recommendations) are given for Low, Medium, and High. (Very High is typically a stretch goal. For Very Low, the supporting explanation in the assessment should identify the architecture's limitations.) The definitions and recommendations are cumulative; for example, recommendations for Medium build on those for Low.

Table 15. Definitions of Assessment Levels and General Recommendations for Adaptive Response

Factor	Level	Definition and General Recommendations
Breadth of response: How many different responsive actions does the architecture support?	Low	**Definition**: The architecture accommodates shutting down of components or communications, as well as restart and/or recovery, and does not preclude reconfiguration. **Recommendations**: Provide guidelines for administrator SOPs and defender TTPs/CCoAs that ensure that response actions do not result in new vulnerabilities (e.g., recovery to a state in which privileges are not properly restricted).
	Medium	**Definition**: The architecture accommodates or includes multiple response actions. **Recommendations**: Provide guidelines for administrator SOPs and defender TTPs/CCoAs to prioritize and make effective use of available mechanisms.
	High	**Definition**: The architecture includes mechanisms for the full range of response actions. **Recommendations**: As part of design and implementation, analyze the interactions of response mechanisms to ensure that they do not interact badly. Provide guidelines for administrator SOPs and defender TTPs/CCoAs to prioritize and make effective use of available mechanisms.
Depth of response: At how many architectural layers, or for how many architectural components, can responsive actions be taken?	Low	**Definition**: The architecture does not preclude response mechanisms for at least one or two layers. **Recommendations**: Identify the layers in the architecture where response mechanisms are not precluded. Identify the costs and benefits associated with implementing different response actions at those locations.
	Medium	**Definition**: The architecture accommodates automated or semi-automated response for at least one or two layers. **Recommendations**: Include requirements for a common or consistent administrative interface that facilitates coordination of response activities at different layers. Provide guidelines for administrator SOPs that ensure that responses taken at different layers do not result in vulnerabilities.
	High	**Definition**: The architecture includes capabilities for semi-automated and, where feasible, fully automated response. **Recommendations**: Ensure that data used to determine and direct response activities are properly protected.
Dynamism: How quickly can response actions be taken?	Low	**Definition**: The architecture requires administrators to direct responsive actions; thus, response speed is a function of operational processes. **Recommendations**: Provide guidelines for administrator SOPs to facilitate rapid response.
	Medium	**Definition**: The architecture accommodates automated or semi-automated response. **Recommendations**: Include requirements for a common or consistent administrative interface that facilitates coordination of response. Include requirements for limited capabilities for coordination of response actions. Provide guidelines for administrator SOPs that ensure that response actions do not create new vulnerabilities.

Factor	Level	Definition and General Recommendations
	High	**Definition**: The architecture includes capabilities for automated and semi-automated response capabilities, including monitoring of the effectiveness and collateral effects of response activities. **Recommendations**: Provide guidelines for administrator SOPs and defender TTPs/CCoAs to make use of available mechanisms, monitor for undesired or unexpected effects, and assess the effectiveness of actions.
Integration: How well are other resiliency technologies are integrated into response?	Low	**Definition**: The architecture provides redundant resources that can be reallocated. Response is informed by Analytic Monitoring, to identify resources suspected of being compromised. **Recommendations**: Provide guidelines for administrator SOPs and defender TTPs/CCoAs to make use of redundant resources while minimizing their potential compromise (e.g., by reconfiguring or locking down redundant resources before reallocating them). Plan to incorporate Diversity together with Redundancy.
	Medium	**Definition**: The architecture provides alternative resources (combining Redundancy with Diversity) that can be reallocated. The architecture supports isolation of resources suspected of being compromised, or of high-value resources to prevent their compromise (applying Segmentation). **Recommendations**: Provide guidelines for administrator SOPs and defender TTPs/CCoAs to make effective use of isolation while minimizing mission impacts.
	High	**Definition**: The architecture includes administrator and defender interfaces to support coordination of reconfiguration, resource reallocation, and dynamic composability across multiple administrative domains (integration with Coordinated Defense), in a manner that takes mission dependencies into consideration (integration with Dynamic Representation). The architecture includes interfaces between response mechanisms and Analytic Monitoring mechanisms, to support automated as well as semi-automated response. Recommendations:

Table 16. Definitions of Assessment Levels and General Recommendations for Analytic Monitoring

Factor	Level	Definition and General Recommendations
Sensor locations[27]: At how many locations is monitoring performed?	Low	**Definition**: The architecture does not preclude monitoring for at least one or two locations. **Recommendations**: Identify the layers or locations in the architecture where monitoring is not performed, but could be. Identify the costs and benefits associated with implementing monitoring at those locations. Plan to acquire and maintain monitoring implementations for at least two layers in the architecture (e.g., network IDS, malware detection).
	Medium	**Definition**: The architecture accommodates or implements monitoring at multiple layers and locations, as well as data fusion and analysis. **Recommendations**: Identify the layers and locations in the architecture where monitoring could be or is performed. Define an approach to implementing monitoring at those layers / locations, and for data fusion and analysis.
	High	**Definition**: The architecture identifies multiple locations at which monitoring is expected or required, and defines information flows and processes for data fusion and analysis. **Recommendations**: Provide guidelines for administrator SOPs and defender TTPs/CCoAs that take advantage of that sensor data fusion and analysis.

[27] Sensors can be located at any of the following places in a network: enterprise perimeter, enclave perimeters, enclave-internal. Similarly, within a system sensors can be located at any of the following places: hypervisor / VMM; OS; distributed application infrastructure; application.

Factor	Level	Definition and General Recommendations
Sensor coordination: How well can sensor coverage and analysis be coordinated within the architecture?	Low	**Definition**: Sensors are configured, and their outputs analyzed, separately. **Recommendations**: Provide guidelines for administrator SOPs that ensure that different sensors are configured in a way that minimizes gaps in coverage. Provide guidelines for administrator SOPs and defender TTPs/CCoAs that involve coordinated analysis of outputs from different sensors.
	Medium	**Definition**: The architecture accommodates automated or semi-automated configuration of different sensors, and limited correlation or data fusion. The architecture accommodates procedural or semi-automated determination of coverage. **Recommendations**: Include requirements for a common or consistent administrative interface that facilitates coordination of sensor coverage. Include requirements for limited capabilities for data fusion or correlation of sensor outputs. Provide guidelines for administrator SOPs that ensure that different sensors are configured in a way that minimizes gaps in coverage. Provide guidelines for administrator SOPs and defender TTPs/CCoAs that involve coordinated analysis of outputs from different sensors.
	High	**Definition**: The architecture includes capabilities for semi-automated and, where feasible, fully automated configuration of different sensors, so that monitoring can be intensified in response to anomalies or I&W. The architecture provides capabilities for data fusion, correlation, and data mining of sensor data. **Recommendations**: Ensure that sensor configuration data and analysis results are properly protected.
Sensor dynamism: How quickly can sensors be recalibrated?	Low	**Definition**: The architecture enables administrators to recalibrate sensors; thus, recalibration speed is a function of operational processes. **Recommendations**: Provide guidelines for administrator SOPs to facilitate rapid recalibration.
	Medium	**Definition**: The architecture accommodates automated or semi-automated sensor recalibration. **Recommendations**: Include requirements for a common or consistent administrative interface that facilitates coordination of sensor coverage. Include requirements for limited capabilities for data fusion or correlation of sensor outputs. Provide guidelines for administrator SOPs that ensure that different sensors are configured in a way that minimizes gaps in coverage. Provide guidelines for administrator SOPs and defender TTPs/CCoAs that involve coordinated analysis of outputs from different sensors.
	High	**Definition**: The architecture includes mechanisms for automated or semi-automated sensor recalibration. **Recommendations**: Include requirements for administrator oversight and override of automated mechanisms. Include requirements to protect the integrity of recalibration mechanisms, and to apply privilege restriction.
Analysis timeliness: How quickly can analysis of sensor or other data be performed?	Low	**Definition**: The architecture supports off-line or asynchronous analysis. **Recommendations**: Provide guidelines for defender TTPs/CCoAs to perform forensic / malware analysis and damage assessment, using available data. Include requirements for protection of monitoring data and captured information.
	Medium	**Definition**: The architecture accommodates limited automated analysis of monitoring data, and includes mechanisms for limited analyst-directed damage assessment. **Recommendations**: Include requirements for administrator and defender interfaces to support efficient analysis. Include requirements for protection of analysis processing and results, and of communications among analysts.
	High	**Definition**: The architecture includes mechanisms for ongoing automated analysis of monitoring data, and for limited ongoing damage assessment and forensic analysis. **Recommendations**: Include requirements for administrator and defender interfaces to visualize the results of damage assessment and forensic analysis.
Scope: What is the scope of analysis?	Low	**Definition**: The architecture supports analysis of monitoring data. (Malware and forensic analysis are assumed to be handled by a separate system.) **Recommendations**: Provide guidelines for administrator SOPs and defender TTPs/CCoAs to use monitoring data and logs perform forensic analysis and damage assessment.

Factor	Level	Definition and General Recommendations
	Medium	**Definition**: The architecture supports semi-automated analysis of monitoring data and damage assessment. (Malware analysis, and forensic analysis beyond basic damage assessment, are assumed to be handled by a separate system.) **Recommendations**: Include requirements to enable analysis to be directed and focused on specific cyber resources or classes of resources. Provide guidelines for administrator SOPs and defender TTPs/CCoAs to aggregate and correlate results of analysis and damage assessment across multiple architectural layers and/or architectural elements (e.g., subnets).
	High	**Definition**: The architecture includes mechanisms for analysis of monitoring data, damage assessment, and forensic analysis. **Recommendations**: Include requirements to isolate forensic and malware analysis from ongoing system/mission operations.

Table 17. Definitions of Assessment Levels and General Recommendations for Coordinated Response[28]

Factor	Level	Definition and General Recommendations
Breadth of defense: How many defensive techniques are applied at a given architectural layer?	Low	**Definition**: One defensive technique (e.g., a standard security control such as access control) is applied. **Recommendations**: In plans of action and milestones (POA&Ms), identify additional defensive techniques that can be applied. Use the results of risk assessments to identify potential attack vectors or adversary tactics, techniques, and procedures (TTPs), and use these to prioritize the additional defensive techniques.
	Medium	**Definition**: Multiple defensive techniques are applied, or the same defensive technique (e.g., malware scanning) is implemented in multiple ways. **Recommendations**: In administrator SOPs and defender TTPs or cyber courses of action (CCoAs), identify dependencies and possible interactions among multiple techniques or implementations, and provide guidance on how to ensure that no conflicts or inconsistencies are introduced.
	High	**Definition**: The architecture includes multiple defensive techniques and, in some cases, multiple implementations of the same technique. **Recommendations**: As part of design and implementation, analyze the interactions of defense mechanisms to ensure that they do not interact poorly. Provide guidelines for administrator SOPs and CCoAs to make effective use of available mechanisms, coordinating the use of different mechanisms at different layers.
Depth of defense: At how many architectural layers, or for how many architectural components, is a given defensive technique applied?	Low	**Definition**: The architecture applies a given defensive technique to a single layer or a single component. **Recommendations**: Perform an architectural analysis to determine whether and how the given technique could be applied at multiple layers. Identify the locations in the architecture where additional implementation of the defensive mechanism is not precluded. Identify the costs and benefits associated with implementing the defensive mechanism at those locations.
	Medium	**Definition**: The architecture applies the given defensive technique to two or more layers or components. **Recommendations**: As part of design and implementation, analyze the applications of the technique to ensure that policies can be enforced consistently. Include guidance on consistent configuration and management in administrator SOPs.
	High	**Definition**: The architecture applies the given defensive technique across multiple contiguous layers or components. **Recommendations**: As part of design and implementation, analyze the interactions of the implementations of the technique at different locations to ensure that policies can be enforced consistently, even while defensive actions are being taken. In POA&Ms, include requirements for a common or consistent administrative interface that facilitates coordination of defense activities at different layers.

[28] Note that the architecture must be described in enough detail that it can be analyzed to identify architectural layers, identify defensive techniques, and perform a mapping from techniques to layers to enable assessment of breadth of defense and depth of defense.

Factor	Level	Definition and General Recommendations
Internal consistency / coordination: How consistently and with how much coordination are cyber defenses, supporting security controls, and supporting performance controls managed within a given administrative span of control (i.e., within a given system, shared service, or common infrastructure)?	Low	**Definition**: Coordination and consistency checking are informal or ad-hoc processes. **Recommendations**: Provide guidelines for administrator SOPs and defender CCoAs/TTPs that identify applicable coordination and information sharing relationships. Perform an architectural analysis to identify a list of critical defensive resources, and how they can be checked for consistent policy enforcement.
	Medium	**Definition**: The architecture incorporates the use of techniques or mechanisms for consistency checking and coordination of diverse mechanisms and/or implementations, within an architectural layer and across contiguous architectural layers or components. **Recommendations**: Provide guidelines for administrator SOPs and defender CCoAs/TTPs to prioritize and make effective use of available mechanisms. Identify the dependencies and interactions among cyber defenses, security controls and performance controls. Provide guidelines for administrator SOPs that ensure that actions taken at contiguous layers or components do not result in vulnerabilities or additional incidents. Identify stakeholders in coordinated defense activities, to ensure that their activities work in concert with each other.
	High	**Definition**: The architecture employs mechanisms for a full range of consistency checking of diverse mechanisms and/or implementations within and across architectural layers, and supports processes for coordination with internal entities (e.g., among system administrators, database administrators, network administrators, cloud service administrators, and cyber defenders), across all layers and components. **Recommendations**: As part of design and implementation, analyze the interactions of coordination support mechanisms to ensure that they do not interact poorly. Provide guidelines for administrator SOPs to prioritize coordination with entities and make effective use of available mechanisms. Ensure that COOP plans include locations and uses of security controls. Ensure that data used to determine and direct defensive activities are properly protected.
External consistency / coordination: How consistently and with how much coordination are cyber defenses, supporting security controls, and supporting performance controls managed across different administrative spans of control?	Low	**Definition**: Information sharing, coordination, consistency checking are informal or ad-hoc processes. **Recommendations**: Provide guidelines for administrator SOPs and defender TTPs/CCoAs that identify applicable coordination and information sharing relationships. Perform architectural analyses of the different systems, shared services, and common infrastructures to identify a list of critical defensive resources, and how they can be checked for consistent policy enforcement.
	Medium	**Definition**: The architecture(s) accommodate or include mechanisms for consistency checking and coordination. Memoranda of Understanding (MOUs), Memoranda of Agreement (MOAs), and/or other instruments are in place to establish responsibilities and expectations. **Recommendations**: Provide guidelines for administrator SOPs and defender TTPs/CCoAs to prioritize and make effective use of available mechanisms and agreements.
	High	**Definition**: The architecture(s) include mechanisms for a full range of consistency checking, and supports processes for coordination with external entities (e.g., among managed service providers for cloud, security, and network services; with cyber defenders for an infrastructure sector). **Recommendations**: Provide guidelines for administrator SOPs and defender CCoAs/TTPs to prioritize coordination with entities and make effective use of available mechanisms. Ensure that COOP plans are coordinated and consistent.

Table 18. Definitions of Assessment Levels and General Recommendations for Diversity

Factor	Level	Definition and General Recommendations
Depth of diversity: At how many architectural layers[29] is diversity provided or supported?	Low	**Definition**: The architecture does not preclude diversity for at least one or two layers. **Recommendations**: Identify the layers in the architecture where diversity is not precluded. Identify the costs and benefits associated with implementing diversity at those locations. Plan to acquire and maintain distinctly different implementations for at least one layer in the architecture.
	Medium	**Definition**: The architecture accommodates diversity at multiple layers. **Recommendations**: Identify the layers in the architecture where diversity is accommodated. Define an approach to implementing diversity at those layers. Plan to acquire and maintain distinctly different implementations at several layers.
	High	**Definition**: The architecture identifies multiple layers at which diversity is expected or required, sets targets for degree of diversity at those layers, and establishes interface standards and guidelines for selecting alternatives, to accommodate diverse technologies. **Recommendations**: Provide guidelines for administrator SOPs that ensure that diversity is maintained. Provide guidelines for defender TTPs that take advantage of that diversity. Note that implementation-specific SOPs & defender TTPs/CCoAs must be aligned or coordinated (see Integration, below).
Breadth of diversity: At how many locations[30] in the architecture is diversity provided or supported?	Low	**Definition**: The architecture does not preclude diversity for at least one or two locations. **Recommendations**: Identify the locations where limited diversity can be expected, and define administrator SOPs & defender TTPs/CCoAs for taking advantage of that diversity. Because compromise of a component can be expected to be replicated quickly to other like components, provide strong monitoring, active defense, and segmentation. Be prepared to isolate components, sub-nets, and systems.
	Medium	**Definition**: The architecture accommodates diversity in multiple locations. **Recommendations**: Identify the locations in the architecture where diversity is accommodated. Define an approach to implementing diversity in those locations. Plan to acquire and maintain distinctly different components / implementations at several locations.
	High	**Definition**: The architecture identifies multiple locations at which diversity is expected or required, sets targets for the number of alternatives at those locations, and establishes interface standards and guidelines for administrator SOPs for selecting alternatives, to accommodate diverse technologies. **Recommendations**: Provide guidelines for administrator SOPs that ensure that diversity is maintained. Provide guidelines for defender TTPs/CCoAs that take advantage of that diversity. Note that implementation-specific SOPs & defender TTPs/CCoAs must be aligned or coordinated (see Integration, below).
Degree of diversity: How many instances / alternatives are expected or accommodated within the selected architectural layers?	Low	**Definition**: Diversity is deprecated in favor of homogeneity. Some limited diversity can be expected by happenstance, e.g., due to presence of legacy components and life-cycle replacement schedules. **Recommendations**: Identify the locations where limited diversity can be expected, and define SOPs for taking advantage of that diversity. Because compromise of a component can be expected to be replicated quickly to other like components, provide strong monitoring, active defense, and segmentation. Be prepared to isolate components, sub-nets, and systems.
	Medium	**Definition**: The architecture accommodates diversity, but does not specify alternative implementations or components. **Recommendations**: When defining an approach to implementing diversity in a location or at a layer, plan to acquire and maintain at least two distinctly different

[29] Examples of system layers include computing platform hardware, VMM / hypervisor, OS, and applications. Examples of network layers include links or paths through the network topology, communications media, encryption devices, firewalls, CDSs, and protocols at different levels in the protocol stack. Examples of data layers include data source, data format, and data transformations.

[30] A location in the architecture roughly corresponds to a box in an architectural diagram.

Factor	Level	Definition and General Recommendations
		components / implementations.
	High	**Definition**: The architecture defines criteria for diversity (i.e., what it means for components to be distinctly different) and sets targets for the number of alternatives at specified locations or layers. **Recommendations**: Ensure that the architectural criteria for diversity take into account technology and supply chain.
Diversity dynamism: How quickly (in terms of technology refreshes or response to incidents or vulnerability discoveries) can new implementations be integrated into the system?	Low	**Definition**: New implementations of components can be made part of future development spirals. **Recommendations**: Ensure that program planning enables new implementations of selected components to be made part of future development spirals.
	Medium	**Definition**: The architecture accommodates new implementations of selected components to be constructed and integrated out-of-cycle with development spirals. **Recommendations**: Ensure that program planning enables new implementations of selected components to be constructed and integrated out-of-cycle with development spirals.
	High	**Definition**: The architecture enables new implementations of selected components to be constructed and integrated in near-real-time. **Recommendations**: Ensure that administrator SOPs and defender TTPs identify the circumstances under which new implementations are to be constructed and integrated.
Integration: How well is diversity integrated with other practices?	Low	**Definition**: The architecture does not preclude consistent privilege restriction or monitoring for alternative implementations. **Recommendations**: Ensure that integration of new implementations includes integration of privilege restriction, consistent across all implementations, and monitoring.
	Medium	**Definition**: The architecture accommodates monitoring and consistent privilege restriction for alternative implementations. **Recommendations**: Ensure that integration of new implementations includes integration of monitoring and consistent privilege restriction capabilities.
	High	**Definition**: The architecture requires monitoring, consistent privilege restriction, and (as feasible) substantiated integrity mechanisms for all alternative implementations. Diversity dynamism is closely integrated with dynamic composability, an Adaptive Response technique. **Recommendations**: Ensure that integration of new implementations includes integration of monitoring, consistent privilege restriction, and (as feasible) substantiated integrity capabilities.

Table 19. Definitions of Assessment Levels and General Recommendations for Privilege Restriction[31]

Factor	Level	Definition and General Recommendations
Depth of Privilege Restriction: At how many architectural layers are privilege restrictions applied?	Low	**Definition**: The architecture enables privilege restriction for at least one or two layers, with emphasis on highly critical and/or sensitive resources. **Recommendations**: Ensure that the administrator SOPs for managing privileges take into consideration the principle of least privilege.
	Medium	**Definition**: The architecture includes privilege restriction at multiple layers. **Recommendations**: Identify the layers in the architecture where privilege restriction in use. Ensure that administrator SOPs call for managing privileges in a consistent and coordinated way at multiple layers.
	High	**Definition**: The architecture includes privilege restriction capabilities at multiple layers, as well as capabilities for coordinated and consistent privilege management. **Recommendations**: Provide guidelines for systems engineering and administrator SOPs that ensure that privileges are managed in a consistent and coordinated way. As part of design and implementation, analyze the interactions of the implementations of the technique at different locations to ensure that privileges can be enforced consistently, even while defensive actions (including failover or recovery as part of COOP) are being taken.
Breadth of Privilege Restriction: How broadly or narrowly is least privilege applied (e.g., is it only applied to services, access to data, individuals)?	Low	**Definition**: The architecture accommodates restrictions of privilege based on a limited number of criteria (e.g., solely based on user identity) associated with resources with a high degree of criticality or sensitivity. **Recommendations**: Ensure that system/program plans provide definitions for privilege restrictions and that administrator SOPs include procedures for restricting access to resources.
	Medium	**Definition**: The architecture accommodates multiple criteria for applying privilege restrictions, and rules for assigning, changing, and removing privileges as well as privilege restrictions on resources. **Recommendations**: Ensure that administrator SOPs for privilege management identify the criteria for privilege restriction, and how those criteria apply to different degrees of resource sensitivity and/or criticality. Provide guidelines for administrator SOPs to ensure that privileges are maintained consistent with established criteria and rules, and are deleted when no longer needed.
	High	**Definition**: The architecture employs multiple criteria for applying privilege restrictions; accommodates rules for assigning, changing, and removing privileges as well as privilege restrictions on resources; and accommodates or provides capabilities for dynamic reassignment of privilege criteria. **Recommendations**: Provide guidelines for systems engineering and administrator SOPs that ensure the criteria for applying privilege restrictions are applied in the implementation and maintained in operational use.
Criticality: To what degree is criticality analysis linked to least privilege?	Low	**Definition**: Privilege restrictions are determined primarily by data sensitivity/criticality. **Recommendations**: Perform (or reuse the results of) a mission / business impact assessment to determine the degree of criticality for resources. Define policies or operational criteria (which can depend on environmental conditions or mission-related situations) for restricting resource use based on criticality as well as sensitivity. Ensure that administrator SOPs for managing privileges are consistent with policies or operational criteria.
	Medium	**Definition**: The architecture accommodates privilege restrictions that reflect resource criticality. **Recommendations**: Ensure that administrator SOPs for privilege management identify how criteria for privilege restriction apply to different degrees of resource criticality.

[31] Note that for an assessment of Privilege Restriction, the architecture must be described in enough detail that resources (including systems, processing, data, and connectivity) can be characterized in terms of degree of criticality and privileges (e.g., criteria for granting access to resources, such as identity, role, or location) can be identified.

Factor	Level	Definition and General Recommendations
	High	**Definition**: The architecture enables privileges and privilege restrictions to change, based on changes to resource criticality. **Recommendations**: Ensure that administrator SOPs for privilege management include procedures for determining whether and how resource criticality is changing (e.g., in response to mission phases), and for checking whether privileges and privilege restrictions are being changed accordingly.
Coordination / consistency: How consistently are privileges defined and assigned? In a system-of-systems, how well are policies and practices coordinated?	Low	**Definition**: Consistency of privileges and privilege restrictions is primarily an artifact of shared identity services. **Recommendations**: Ensure that policies and/or operational criteria for assignment of privileges and privilege restrictions are defined and documented. Provide guidelines for administrator SOPs to ensure that privilege restriction is consistent across all segments / enclaves and defensive actions. Provide guidelines for investment planning to provide privilege restrictions for additional resources (e.g., moderate as well as high criticality and/or sensitivity).
	Medium	**Definition**: The architecture accommodates limited integration of privilege management capabilities across segments / enclaves / systems, and across architectural layers. **Recommendations**: Ensure that systems engineering for restricting resource use based on privileges, including integration of new (and possibly diverse) criteria, will be consistent across different architectural layers or classes of resources. Provide guidance for administrator SOPs to check that privileges are assigned consistent with policy / operational criteria. Ensure that administrator SOPs identify entities with which privilege management must or should be coordinated, and define coordination procedures.
	High	**Definition**: The architecture enables integration of privilege management capabilities across segments / enclaves / systems, and across architectural layers, including consistency checking as well as assessment of how well the assignment of privileges and privilege restrictions matches established policies or operational criteria. **Recommendations**: Provide guidance for administrator SOPs and defender CCoAs/TTPs to ensure that use of redundant resources in CCoAs takes advantage of isolation to ensure privilege are not widely allocated. Provide guidance for administrator SOPs to validate consistency checking, and to document and resolve conflicts.

Table 20. Definitions of Assessment Levels and General Recommendations for Redundancy

Factor	Level	Definition and General Recommendations
Depth of redundancy: At how many layers is redundancy provided?	Low	**Definition**: The architecture accommodates redundancy (as part of backup and restore functionality or basic network topology) for at least one or two layers. **Recommendations**: Define an approach to implementing redundancy. Ensure that the redundancy the architecture provides, and the administrator SOPs for using that redundancy, take into consideration the need to be resilient against cyber threats.
	Medium	**Definition**: The architecture accommodates redundancy at multiple layers. **Recommendations**: Identify the layers in the architecture where redundancy is accommodated. Define an approach to implementing redundancy in a consistent and coordinated way at multiple layers.
	High	**Definition**: The architecture identifies multiple layers at which redundancy is expected or required, and sets targets for validating the redundancy at those layers. **Recommendations**: Provide guidelines for systems engineering and administrator SOPs that ensure that validation of redundancy is maintained. Provide guidelines for defender TTPs that take advantage of that redundancy. Note that implementation-specific SOPs & defender TTPs must be aligned or coordinated (see Integration, below).
Breadth of redundancy: How many duplicate	Low	**Definition**: The architecture accommodates duplication of selected resources. **Recommendations**: Ensure that system/program plans provide local spare copies and/or extra capacity, and that administrator SOPs include procedures for using redundancy.

Factor	Level	Definition and General Recommendations
copies of a given resource exist? Where?	Medium	**Definition**: The architecture accommodates redundancy using multiple locations (e.g., offsite backup) and/or alternative communications paths. **Recommendations**: Ensure that administrator SOPs for all locations are consistent with respect to protecting resources and ensuring their consistency.
	High	**Definition**: The architecture includes hot backups, with the ability to revert to previously stored versions, and sets targets for the number of duplicate copies, spare capacity, and/or alternative communications paths. **Recommendations**: Provide guidelines for systems engineering and administrator SOPs that ensure that validation of redundancy is maintained. Provide guidelines for defender TTPs that take advantage of that redundancy. Note that implementation-specific SOPs & defender TTPs must be aligned or coordinated (see Integration, below).
Validation: How consistent and independent are duplicate copies?	Low	**Definition**: Redundancy is deprecated in favor of cost containment. Some limited redundancy is provided by backup and restore capabilities. **Recommendations**: Provide guidelines for investment planning to provide redundancy beyond backup and restore.
	Medium	**Definition**: The architecture accommodates redundancy, but does not include validation mechanisms. **Recommendations**: Ensure that systems engineering includes analysis of redundant services and communications to identify and mitigate single points of failure. Provide guidelines for administrator SOPs to ensure that the patch status and configuration of software is consistent across duplicate copies of software.
	High	**Definition**: The architecture provides validation mechanisms for redundant resources (e.g., patch status and configuration for software, consistency checking across duplicate information stores). **Recommendations**: Ensure that administrator SOPs and defender TTPs include procedures for responding to detection of inconsistencies.
Integration: How well is redundancy integrated with other practices?	Low	**Definition**: Redundancy is viewed as part of performance engineering and contingency planning. **Recommendations**: Provide guidelines for administrator SOPs to ensure that privilege restriction and analytic monitoring are consistent across all copies.
	Medium	**Definition**: The architecture accommodates limited integration of diversity with redundancy. **Recommendations**: Ensure that integration of new (and possibly diverse) copies includes integration of monitoring and consistent privilege restriction capabilities.
	High	**Definition**: The architecture includes segmentation, so that uncompromised duplicate copies can be isolated from compromised copies, and substantiated integrity mechanisms that can be applied to duplicate copies. **Recommendations**: Provide guidance for administrator SOPs and defender TTPs to ensure that use of redundant resources in CCoAs takes advantage of isolation and substantiated integrity mechanisms.

C.3 Example of an Assessment Scale for Levels of Implementation

Table 21. Levels of Implementation for Diversity

Level	Description	Diversity Implementation
Very High	Fully, effectively, evolvably based on ongoing assessment	High, plus: A documented strategy for maintaining diversity identifies • how diversity will be maintained during Operations and Maintenance (O&M) or as part of acquisition spirals, and • how different components must evolve to maintain diversity in response to diversity-reducing changes in the supply chain (e.g., vendor consolidation).
High	Fully, effectively	A documented architectural diversity strategy identifies • architectural layers at which diversity is implemented, • the target and minimum degree of diversity at each layer, and • interface specifications to ensure that multiple implementations at those layers can be acquired or maintained. The degree of diversity at each layer is evaluated at key milestones in the SDLC, and exceeds the minimum. The CONOPS and SOPs identify • how diverse implementations are configured to ensure consistent enforcement of security policies, and • under what circumstances and how to switch from one implementation to another.
Medium	Partially	The architecture identifies • architectural layers at which diversity will be implemented or accommodated, typically as part of incidental diversity, • interface specifications that enable multiple implementations at those layers to be acquired or maintained, and • a roadmap for phasing in multiple implementations.
Low	Planned	The architecture identifies architectural layers at which diversity could be implemented or accommodated, typically as part of incidental diversity.
Very Low	Not planned	The architecture explicitly eschews diversity (typically to avoid complexity). (Note that this architectural decision entails programmatic risks, if specific components become unavailable or are known to be compromised.)

C.4 Examples of Assessment Scales for Resiliency-Related Cyber Defender Activities

Table 22. Examples of Value Scales for Cyber Resiliency-Enhancing Activities

Activity	Low	Medium	High
Identify, and maintain a representation of, functional and mission dependencies among cyber resources [Dynamic Representation]	Static, manually-generated representation, typically as part of design documentation	Combination of static and dynamic representations, supported by automation (e.g., network maps updated by network mapping services)	Dynamic representation of dependencies, automatically updated based on observation and analysis of usage
Identify mission / business function dependencies on cyber resources [Dynamic Representation]	Static, manually-generated representation, typically as part of design documentation	Combination of static and dynamic representations, supported by automation (e.g., Business Service Management (BSM) or IT Service Management (ITSM) tools)	Dynamic representation of dependencies, automatically updated based on observation and analysis of usage
Identify non-mission / business function dependencies on or uses of cyber resources [Dynamic Representation]	Static, manually-generated representation, typically as part of design documentation	Combination of static and dynamic representations, supported by automation (e.g., Business Service Management (BSM) or IT Service Management (ITSM) tools)	Dynamic representation of dependencies, automatically updated based on observation and analysis of usage
Identify dependencies and interactions among cyber defenses, security controls, and performance controls [Coordinated Defense]	Ad-hoc identification, based on defender and administrator experience; typically undocumented or documented in informal (and unshared) work notes	Semi-structured identification, typically in design documentation, with limited representation in SOPs; documentation augmented by defender and administrator work notes and checklists	Structured identification, based on design documentation, defender and administrator experience, and engineering analysis, reflected in SOPs, configuration management guidance, and MIA/BIA/CJA
Validate data provenance [Substantiated Integrity]	Ad-hoc determination of data provenance, based on manual capture and/or use of data gathered by existing mechanisms (e.g., system logs)	Limited automation for determination of data provenance (e.g., provenance mechanisms integrated with some applications)	Architectural integration of data provenance capture and delivery mechanisms
Validate data integrity / quality to ensure it has not been corrupted [Substantiated Integrity]	Ad-hoc determination of data quality; inferential determination of non-corruption (e.g., manual consistency checks)	Automated determination of non-corruption (e.g., using cryptographic checksums); determination of data quality decoupled from determination of non-corruption	Architectural integration of data quality assessment mechanisms, including determination that potential corruption has not occurred as a dimension of data quality

Activity	Low	Medium	High
Identify key locations to place mechanisms [Coordinated Defense]	Locations for mechanisms are selected based on common practice	Locations for mechanisms selected based on asset / resource criticality	Locations for mechanisms based primarily on functional dependencies and secondarily on asset / resource criticality
Integrate cyber resiliency strategy with other organizational strategies [Coordinated Defense]	Coordination of information security with business continuity; information security is part of larger-scale risk management (e.g., coordinated management of information, IT, compliance, and business risks). (Cyber Prep Level 2)	Consistency between cyber security, architectural, and acquisition strategies; cyber security (including cyber resiliency) is part of enterprise risk management. (Cyber Prep Level 3)	Coordination of architectural and acquisition strategies with cyber security strategy; cyber resilience is a central part of mission assurance strategy, which is part of the organization's mission and enterprise risk management strategies. *or* Full integration of cyber security and resiliency into the organization's mission assurance strategy, which is a significant part of the organization's mission and enterprise risk management strategies. (Cyber Prep Levels 4-5)

Appendix D Cyber Resiliency Techniques

This Appendix provides more detailed discussion of the cyber resiliency techniques presented in [6] as basis for engineering analysis and resiliency roadmap development. For each technique, a basic description presents key concepts and a few references, primarily to recent research. For a survey of the cyber resiliency research landscape through early 2011, see [9]. For some cyber resiliency techniques, a few more specific techniques or classes of technology are described.

To support the development of recommendations, most[32] cyber resiliency techniques are mapped to application domains as described in Section 2.1. The maturity of the technique (or the more specific techniques, if presented) is assessed using Table 23 below.[33]

Table 23. Relative Maturity Levels

Relative Maturity	Description
Highly Mature 5	The technology is available commercially or as GOTS. The technology is in common use. Standards of good practice for its use, based on extensive experience, have been documented. At least some of the "best practices" consider the need for resilience in the face of cyber threats.
Mature 4	The technology is available commercially or as GOTS. Operational experience and guidance have been documented, including some guidance related to resilience in the face of cyber threats. (Corresponds to TRL 8-9)
Transitional 3	Prototype or proof-of-concept technology is integrated into a representative demonstration / experimental environment or is in limited / experimental operational use. (Corresponds to TRL 6-7)
Immature 2	Prototype or proof-of-concept technology has been developed. (Corresponds to TRL 3-5)
Highly Immature 1	Key concepts and approaches are being explored or developed. (Corresponds to TRL 1-2)

D.1 Adaptive Response

Adaptive Response techniques enable systems and organizations 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 course of action (CCoA) that best changes the attack surface, maintains critical capabilities, and restores functional capabilities. Indications that an attack is underway include detection of divergence from the organization's established conditions of normal operations, as well as externally provided threat intelligence. Responses to the attack include changes to the capabilities, processes, technologies, or security postures that were previously presented to the adversary. Examples include employing applications not previously presented to the adversary, changing resource allocations within networks or computing environments, and changing the configuration of networks, systems, or applications.

Adaptive Response includes a mixture of human and automated decisions. Policy- and risk-driven automation will enable systems to evolve toward greater autonomic decision-making.

[32] Some cyber resiliency techniques (Coordinated Defense, Dynamic Representation, Unpredictability) involve integration across multiple layers. For these techniques, no mapping to application domains is performed.

[33] When no evidence for the use of a technique at a given layer can be found, or when the technique is simply not applicable at a layer, the corresponding table cell is left blank.

D.1.1 Existing Techniques

Existing techniques include administrator-directed reconfiguration and resource reallocation, which can include resource reprovisioning (i.e., changing the software, configuration, and data associated with the resource, so it is ready to perform the functions to which it has been reallocated). Configuration changes can affect connectivity or functional dependencies; authorizations; and performance settings. Reconfiguration is supported by management (and security management) tools at multiple architectural layers (e.g., network, operating system, application). Resource reallocation similarly is supported by management tools at multiple layers (e.g., network, operating system, cloud computing infrastructure, application).

D.1.2 Emerging Techniques

Emerging techniques enable CCoAs to be executed without taking components off-line. This minimizes interruptions in mission capabilities.

D.1.2.1 Dynamic Reconfiguration

Dynamic reconfiguration means making configuration changes to a component while it continues operating, as opposed to taking the component off-line or out of service.

Administrator-directed dynamic reconfiguration of some system components is currently feasible, e.g., for firewalls [10]. However, automated dynamic reconfiguration is still in the R&D stages for hardware (chip multiprocessors (CMPs) [11], multiprocessor system-on-a-chip (MPSoC) architectures [12] and for field programmable gate arrays (FPGAs) [14]), firewalls [15], networks [15] (purely from a performance perspective), and IDS (Rehak et al., 2009). Whether reconfiguration is dynamic or not, security of the reconfiguration process presents challenges [16].

Automated dynamic reconfiguration is a central feature of adaptive software, which reconfigures itself based on monitoring [18]. Dynamic reconfiguration can be viewed as a form of moving target defense, and can make use of genetic algorithms to find alternative configurations [18].

Dynamic reconfiguration raises concerns for stability, particularly when a failure occurs during reconfiguration. Operational guidance needs to take rollback (recovery to a known good state) into consideration. Dynamic reconfiguration for systems-of-systems presents additional concerns, related to service level agreements and to the potential for propagating attacks and other faults; technical approaches are being investigated [20].

D.1.2.2 Dynamic Resource Allocation

Dynamic resource allocation/reallocation (i.e., making changes in the allocation of resources to tasks or functions without terminating functions or processes) is typically considered from the viewpoint of performance management. As a general problem, dynamic resource allocation is a challenge in data centers [20] and virtual environments [22] and can be handled via approximation algorithms [22]. A growing body of work applies to cloud computing or other service-oriented middleware, where secure provisioning is an active area of investigation [23]. Products and tools are available for dynamic resource allocation in networks (see, for example, [24]). From the standpoint of security, dynamic resource allocation has been studied in the context of grid computing [26].

Dynamic resource allocation raises issues of adherence to service level agreements (SLAs), including security as a service.

D.1.2.3 Dynamic Composability

Dynamic composability involves dynamic replacement of software components with equivalent functionality performed by different software [26]. Replacement requires both construction of new components and dynamic composition, i.e., integration and optimization of new components into an existing system. Dynamic composition can also be part of system integration [27], and dynamic composable computing (DCC) is of particular importance for mobile platforms [28]. For cyber resiliency, dynamic composability applies primarily to mission/business function applications and services, where it is currently immature. Dynamic composability can also apply to other software layers. Dynamic composability changes the attack surface, forcing the adversary to develop new or adapt existing of malware or attacks.

D.1.2.4 Proactive Recovery/Proactive Resilience

Proactive recovery techniques were initially explored in the context of fault-tolerant systems and networks [29] [30], and subsequently for intrusion-tolerant systems. Issues with proactive recovery in networks have been explored [31], and resolution of these issues explored as proactive resilience [32]. Proactive recovery techniques are also relevant in virtual environments [33] [34], where they may be more appropriately considered a form of Non-Persistence.

D.1.3 Applicability and Maturity

The following table presents an overall assessment of the maturity of Adaptive Response techniques and technologies, as applied to different architectural layers[34].

Table 24. Applicability and Maturity for Adaptive Response Techniques

Application Domain	Dynamic Reconfiguration	Dynamic Resource Reallocation	Dynamic Composability
Hardware/firmware	Immature.		
Networking/ communications	Mature; Transitional in the cyber resilience context.	Mature; Transitional in the cyber resilience context.	Highly Immature.
System/ network component	Transitional.		
Operating system	Mature; Transitional in the cyber resilience context.		Highly Immature.
Cloud, virtualization, and/or middleware infrastructure	Mature; Transitional in the cyber resilience context.	Mature; Transitional in the cyber resilience context.	Highly Immature.
Mission / business function application / service	Dependent on application/service.	Dependent on application/service.	Immature.
Information stores	Mature in the context of database tuning. Immature in the cyber resilience context.		
Information streams / feeds	Immature, particularly in the cyber resilience context; current practice relies on manual reconfiguration.		
System / system-of-systems	Immature, particularly in the cyber resilience context.	Immature, particularly in the cyber resilience context.	

[34] Proactive recovery applies primarily to networking, and is immature.

D.2 Analytic Monitoring

Analytic Monitoring techniques *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 (e.g., a mission/business segment, a common infrastructure, a set of shared services, or a system). 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 post-mortem), and use of validation techniques to identify changes in infrastructure that indicate an ongoing attack.

D.2.1 Existing Techniques

To *identify potential adversary activities*, existing techniques include security monitoring (particularly as performed by intrusion detection systems (IDS)), performance monitoring, attack sensing and monitoring (AS&W), and cyber situational awareness [35] applications. Analytic Monitoring differs from security and performance monitoring in its emphasis on informing defender actions by

- Finding indications of a stealthy and well-resourced adversary;
- Detecting and assessing damage; and
- Watching for adversary activities during recovery and evolution.

Penetration testing or active probing can be used to determine IDS effectiveness. Techniques and tools for malware and forensic analysis (i.e., analysis of artifacts left behind by adversary activities) and damage assessment (i.e., analysis of behavior, data, and system artifacts to determine the presence and extent of damage) are used to *analyze adversary activities*; a variety of free and commercial tools are available. Analysis can drive dynamic sensor reconfiguration.

Analytic Monitoring can make use of information sharing – e.g., about attack trends, vulnerabilities, and the results of forensic analysis – with other organizations [7]. Structuring of threat data facilitates information sharing [37] [2].

D.2.2 Emerging Techniques

Monitoring – ongoing collection and preliminary analysis of data – can be problematic in virtualized or cloud computing environments [38], where the tracking of state information within a virtual machine can disappear when the VM goes away. Some VM monitoring can be integrated with security monitoring, and security monitoring can be performed at the hardware layer [39] and can be integrated across different layers in a cloud environment [40] [41]. Monitoring can also be performed by the hypervisor [42].

Monitoring can also be problematic when the devices where monitoring must be performed have power or bandwidth limitations (e.g., mobile devices, wireless sensor networks). In such environments, Analytic Monitoring must be integrated with Adaptive Response, so that monitoring adapts to environmental challenges and constraints [43] [44].

Intrusion detection for mobile ad-hoc networks remains a research challenge [45]. Network analytic techniques can be extended to consider the presence of compromised nodes [46]. Intrusion detection techniques historically have looked for atomic events; correlation and fusion for multi-stage attacks remains a research topic [47].

Cyber situational awareness, including cyber sensor fusion and analysis (i.e., correlation and aggregation of monitoring data from multiple sources, and near-real-time analysis), is an active research area [35]. So-called "big data analytics" are being integrated into network monitoring tools [48] and are being applied to operating systems in large-scale environments [49].

Forensic techniques and tools for information stores remain investigatory [50] [51], despite some established practices and tools. Forensic analysis remains labor-intensive, and can benefit from tools and processes for automation [52]. Cyber damage assessment remains investigatory [53] [54], but can be facilitated by big data analytics [55] and Substantiated Integrity mechanisms [56].

D.2.3 Applicability and Maturity

The following table presents an overall assessment of the maturity of Analytic Monitoring techniques and technologies, as applied to different architectural layers.

Table 25. Applicability and Maturity for Analytic Monitoring Techniques

Application Domain	Monitoring	Sensor Fusion and Analysis	Malware and Forensic Analysis
Hardware/firmware	Immature.	Immature.	Immature (largely rely on tamper-evident technology).
Networking/ communications	Mature for conventional IDS; Transitional in the cyber resilience context or for MANETs; Transitional-to-Mature for big data analytics.	Mature for conventional enterprise networks; Immature-to-Transitional for analysis of multi-stage events.	Mature (network forensics).
System/ network component	Mature for some components; immature or transitional for others.	Immature.	Mature for some (e.g., mobile device forensics); immature or transitional for others.
Operating system	Mature; Transitional-to-Mature for big data analytics.	Mature.	Mature (focus of most malware analysis tools).
Cloud, virtualization, and/or middleware infrastructure	Immature.	Immature.	Immature.
Mission / business function application / service	Dependent on application/service.	Dependent on application/service.	Mature for malware analysis; immature for forensic analysis.
Information stores	Mature for DBMSs. Immature for unstructured information stores.		Immature-to-Transitional for DBMSs. Immature for unstructured information stores.
Information streams / feeds	Immature.		
System / system-of-systems	Immature-to-Transitional.	Immature-to-Transitional.	Immature.

D.3 Coordinated Defense

Coordinated Defense techniques *manage adaptively and in a coordinated way multiple, distinct mechanisms to defend critical resources against adversary activities*. Coordinated Defense relies

on technical defense-in-depth: using multiple mechanisms to apply the same technique to different technologies or architectural layers, and using distinct mechanisms to apply different techniques. Greater asset criticality merits greater layering. 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. Defense-in-depth is well understood and accepted in principle. In practice, issues arise related to governance, as well as outsourcing and partnership relationships which can limit visibility into which or how many mechanisms are used.

Adaptive management entails changing how defensive mechanisms are used (e.g., making configuration changes, turning on some mechanisms while turning off others, deciding when and how to update or patch software) based on changes in the operational environment (e.g., changes in mission/business needs or priorities) as well as changes in the threat environment (e.g., notification of newly discovered vulnerabilities in component technologies). Adaptive management requires the ability to visualize the consequences of administrator action on missions as well as on cyber resources. Thus, Coordinated Defense depends on visualization and analysis techniques that also support Dynamic Representation, Analytic Monitoring, and Adaptive Response. A significant issue is how to provide staff with the information they need to make decisions that fall within their authority – and only that information, to avoid information overload. Resolution of that issue involves governance as well as technical solutions [7], and typically relies on a multi-tiered approach to risk management [57] [58].

Cyber defenses, supporting security controls, and supporting performance controls must be managed in a consistent manner across multiple administrative spans of control. *Coordination and consistency analysis* – determining operational consequences of changes and ensuring that changes do not create inconsistent defenses – is essential to ensure that an attack that involves one defensive mechanism does not create adverse unintended consequences (e.g., lockout, cascading alarms) by interfering with another defensive mechanism. Coordinated cyber defenses must take dependencies into consideration [59] [60]. Thus, cyber courses of action (CCoAs) and contingency plans must be defined in a coordinated way.

Organizations develop administrator standard operating procedures (SOPs), as well as contingency and continuity of operations plans. Increasingly, cyber defense considerations are represented in these, as well as in cyber defense playbooks and CCoAs. Coordinated management is facilitated by enterprise security management tools and product suites. Research and development activities are experimenting with and transitioning cyber operations management capabilities [61].

The effectiveness of Coordinated Defense is enhanced when combined with such other techniques. Coordinated Defense involves coordination of mechanisms across architectural layers; therefore, no discussion of applicability to different layers is meaningful.

Table 26. Maturity and Related Techniques for Coordinated Defense

	Technical Defense-in-Depth	Coordination & Consistency	Adaptive Management
Maturity	Mature, with some aspects Transitional for emerging technologies.	Transitional-to-Mature, depending on governance and interoperability.	Transitional.
Related techniques	Diversity Dynamic Representation Segmentation	Analytic Monitoring (coordinated placement of sensors) Privilege Restriction	Adaptive Response

D.4 Deception

Deception techniques *use obfuscation and misdirection (e.g., disinformation) to confuse or mislead an adversary*. Deception can make the adversary uncertain how to proceed, delay the effect of the adversary's attack,[35] increase the risk to the adversary of being discovered, or expose an adversary's tradecraft. Deception can take the form of dissimulation ("hiding the real") or simulation ("showing the false").[36]

D.4.1 Obfuscation and Dissimulation

Obfuscation (or dissimulation) techniques include masking (e.g., using encryption or function hiding), repackaging (e.g., using data transformation), and dazzling (e.g., responding to adversary queries with deliberately confusing or erroneous information). Encryption is widely used, particularly in networking and data storage. Encrypted processing is feasible for some applications [62]. Self-encrypting drives (SEDs) are commercially available, consistent with the Trusted Computing Group's specification. Hardware obfuscation techniques have been developed, primarily for protection of intellectual property but also for security [63].

D.4.2 Misdirection and Simulation

Misdirection and simulation techniques include inventing (e.g., simulating a non-existent application), mimicking (e.g., fabricating documents [64] or data stores), and decoying (e.g., using honeypots). Honeypot and honeynet technologies have been made available via a variety of initiatives, including the open source Honeynet project (http://www.honeynet.org/), the UK Honeynet project (http://www.ukhoneynet.org/), and the Web Application Security Consortium project. While honeypots and honeynets are increasingly recommended, standards of practice for implementing and using them have not been promulgated. Large-scale deception environments present a research challenge [65].

D.4.3 Applicability and Maturity

The following table presents an overall assessment of the maturity of Deception techniques and technologies, as applied to different architectural layers.

[35] Potential benefits of delaying the attack can include providing the organization additional time to complete critical mission functions, as well as providing time to deploy an adaptive response.
[36] See [153], cited in [165], [164].

Table 27. Applicability and Maturity for Deception Techniques

Application Domain	Obfuscation	Simulation
Hardware/firmware	Mature for self-encrypting drives.	Immature.
Networking/ communications	Mature (using encryption, obfuscation of IP addresses).	Transitional.
System/ network component	Mature for self-encrypting drives. Mature software-based encryption for mobile devices; however, encryption coverage may be incomplete.	
Operating system	Mature (to avoid OS "fingerprinting"). (Note that this is a known adversary technique.)	Transitional.
Cloud, virtualization, and/or middleware infrastructure	Immature.	Transitional.
Mission / business function application / service	Dependent on application/service.	Dependent on application/service.
Information stores	Mature for stored data (using encryption). Immature-to-Transitional for encrypted query processing.	
Information streams / feeds	Mature (using encryption).	
System / system-of-systems		Transitional (honeypots, honeynets).

D.5 Diversity

Diversity techniques *use a heterogeneous set of technologies, communications paths, suppliers, and data sources to minimize the impact of attacks and force adversaries to attack multiple different types of technologies.* Technologies include hardware, software, firmware, and protocols. One mechanism for implementing diversity for software is virtualization, which allows rapid, inexpensive changes in applications, thus making some forms of diversity easier to implement.

Diversity requires that technologies that provide the same (or equivalent) functionality differ enough that they do not present the same attack surface to an adversary. Examples of methods to determine whether two instances are different include data pedigree, functional dependency analysis, hardware or software component pedigree as established by supply chain risk management (SCRM), and use of alternative specifications for automatically generated software.

Diversity is a commonly cited technique for resilience [66] [67]. Another term for diversity is heterogeneity [68]. As noted in [69], Diversity is vital to effective Redundancy. Therefore, Diversity and Redundancy are often analyzed together [70].

D.5.1 Existing Techniques and Technologies

For information and communications technology, existing techniques include architectural diversity and design diversity.[37]

[37] For safety-critical systems, six categories of diversity have been defined: design diversity, equipment diversity, functional diversity, human diversity, signal diversity, and software diversity [165]. In this document, "design diversity" includes design, functional, human, and software diversity; "data diversity" includes signal diversity; and hardware diversity is viewed as an application of design diversity to hardware and firmware.

D.5.1.1 Architectural Diversity

Architectural diversity is the accommodation in an architecture for different components that provide the same functionality (e.g., OSs, VMMs, servers). Architectural diversity can be *planned*. For example, a reference architecture can explicitly identify alternative components, or can identify technical standards that focus on behavior and interface specifications. For networking, architectural diversity can involve the use of multiple protocols, multiple communications media (e.g., satellites, wireless, land lines), and multiple communications paths. Architectural diversity for networking thus can be used in conjunction with Segmentation.

Alternately, architectural diversity can be *incidental*, a result of decisions that are not primarily architectural in nature. Some diversity occurs in enterprise systems due to the presence of multiple instances (e.g., versions, configurations) of the same products, typically acquired at different times and/or by different organizational units. Although this diversity is incidental to the enterprise architecture (and may, indeed, be viewed as undesirable from a management viewpoint), it can provide the benefit of improving attack detection [71].

Increasingly, the trend toward bring-your-own-device (BYOD) leads to diversity in end-user devices (mobile systems and network components) [72]. Thus, a "managed diversity" approach to enterprise architecture has been recommended [73] [72].

D.5.1.2 Design Diversity

Design diversity, in which different designs (and subsequently different implementations) are developed based on the same requirements, is an established concept for fault-tolerance. For software, the best-known technique is N-version programming, in which multiple implementations are created by different programming teams. Design diversity demonstrates clear benefits [74]. However, cost is a significant consideration, particularly for hardware [75]. Therefore, design diversity is applied primarily to safety-critical systems.

While considerable research was performed in the 1990s in the area of metrics and techniques for assessing software diversity, practical applications have lagged [76]. Thus, the research area of metrics and modeling to assess software diversity remains active [77].

D.5.2 Emerging Techniques and Technologies

Research to support cost-benefit analysis has so far been limited [78]. The question of whether two implementations truly provide the same functionality is difficult to answer without extensive testing, and the question of how different two implementations are presents its own challenges. In particular, different components or applications, offered by different vendors, can incorporate common hardware, firmware, or software libraries. Thus, diversity for COTS components must be aligned with supply chain risk management [79].

D.5.2.1 Implementation or Synthetic Diversity

For software, implementation or synthetic diversity [80] involves transforming implementations, using such techniques as instruction set randomization (ISR), address space randomization (ASR), and data space randomization (DSR) [81]. This use of memory-based diversity is sometimes referred to as K-variants [82]. Similarly, N-variant systems generate variants of implementable software, and can be used with N-version programming [83]. For applications distributed to mobile devices, a large-scale combination of randomization and implementation diversity might be considered [84].

Virtualization enables software diversity to be applied in an operational environment [85]. Some dynamic positioning or moving target techniques can also be characterized as artificial diversity [86].

D.5.2.2 Information Diversity

Information diversity[38] can be synthetic or inherent. Inherent information diversity uses different data sources; the determination of whether and how different data sources can be used is highly mission-dependent, and creates challenges for integration and analysis [87]. Tracking of provenance and pedigree could enable users to determine whether and how diverse the data actually is; many different approaches to data provenance[39] have been identified [88].

With synthetic diversity, different variants of the same information are generated automatically [89]. Randomization of parameters can also provide diversity [90].

D.5.3 Applicability and Maturity

The following table presents an overall assessment of the maturity of diversity techniques, as applied to different architectural layers.

Table 28. Maturity of Diversity Techniques

Application Domain	Specific Technique(s)	Assessment
Hardware/firmware	Design diversity	Mature, but costly; need to consider SCRM
Networking/communications	Design diversity Alternate communications paths and media (see Segmentation)	Mature for wireless communications (used to improve performance) [91] Mature but costly for different communications media
System/network component	Design diversity Implementation diversity	Mature, but often more apparent than real Mature, but often more apparent than real; need to consider SCRM
Mobile system/network component	Design diversity Implementation diversity	Immature – diversity is a consequence of marketplace in flux Transitional
Operating system	Design diversity Implementation diversity	Mature, but limited Transitional
Cloud, virtualization, and/or middleware infrastructure	Design diversity Implementation diversity	Immature – diversity is a consequence of marketplace in flux
Mission / business function application / service	Design diversity Implementation diversity	Immature-to-Transitional
Information stores	Information diversity	Depends on diversity of information streams/feeds
Information streams / feeds	Information diversity	Immature-to-Transitional

[38] The term “information diversity” is used to avoid confusion with “data diversity” and “informational diversity.” Data diversity refers to a specific approach to fault-tolerance, generating and executing a set of automatically diversified variants on the same inputs. Data diversity can be used in N-variant systems [89]. Informational diversity refers to differences in educational background and experience among members of a team [145].

[39] Data provenance is also a mechanism for Substantiated Integrity, and is discussed in greater detail in Section D.13.

D.6 Dynamic Positioning

Dynamic Positioning techniques *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. As with Coordinated Defense, this increases the chance of adversaries revealing their actions and tradecraft. Dynamic Positioning applied to sensors supports Analytic Monitoring by allowing the monitoring of activities in specific parts of a system or involving specific assets to be adjusted in consideration of threat, vulnerability, or anomaly information. Examples of technologies to support this technique include virtualization and distributed processing.

D.6.1 Existing Technologies

Existing middleware infrastructures enable administrators to use distributed processing, and to allocate resources in such a way that services and information assets will be relocated. Distributed database technology is also mature. For communications, dynamic positioning techniques include frequency hopping and mechanisms for rotating or changing IP or MAC addresses.

D.6.2 Emerging Technologies

Considerable research is ongoing in the area of moving target defenses [92] [93] [94], and for supporting mechanisms such as state snapshotting [95]. While the phrase "moving target" suggests relocation [96] or changes in networking [97] [98], numerous other topics are often identified as part of moving target defenses. The following table shows how the areas identified by the National Symposium on Moving Target Research [92] are covered in this document. Dynamic relocation involves state capture and state restoration on a different platform, enabled by platform or OS independence [96] or by virtualization [99]. IPv6 offers additional opportunities for network repositioning [100].

Table 29. Topics from Moving Target Research Symposium

Moving Target Topic	Topic or Cyber Resiliency Technique
Dynamic network services	Part of Adaptive Response
Game theoretic approaches	Can inform development of CCoAs and strategies for Coordinated Defense and Adaptive Response.
Virtual machines	Vital for Dynamic Positioning and Non-Persistence. Constitute a challenge for Analytic Monitoring and Dynamic Representation.
Cloud computing	Constitutes an application domain for resiliency techniques.
Dynamic execution	Part of Adaptive Response
Automated response actions	Part of Adaptive Response
Situational awareness	Intersection of Coordinated Defense and Analytic Monitoring
Artificial diversity	Part of Diversity
Encryption to dynamically hide network and transport layer addresses	Example of Deception
Dynamic reconfiguration	Part of Adaptive Response

D.6.3 Applicability and Maturity

The following table presents an overall assessment of the maturity of dynamic positioning techniques, as applied to different architectural layers.

Table 30. Applicability and Maturity for Dynamic Positioning Techniques

Application Domain	Distributing Assets	Repositioning Assets
Hardware/firmware, System/ network component	Mature (physical distribution across multiple facilities).	Mature for some missions (physical relocation in tactical environments).
Networking / communications	Mature for rotating IP or MAC addresses.	Mature for frequency hopping. Mature for changing IP or MAC addresses; Immature-to-Transitional for IPv6.
Cloud, virtualization, and/or middleware infrastructure	Mature (service oriented architecture (SOA) middleware).	Mature as an enabler for repositioning (using non-persistence and resource reallocation).
Mission / business function application / service	Dependent on whether application/ service has been designed for SOA.	Immature-to-Transitional.
Information stores	Mature for distributed databases.	Immature.

D.7 Dynamic Representation

Dynamic Representation techniques *construct and maintain dynamic representations of components, systems, services, mission dependencies, adversary activities, and effects of alternative cyber courses of action*. A representation is *dynamic* if it can reflect changes in state or behavior. A static representation (e.g., a network diagram that does not allow for differences in mission criticality of network components depending on which mission functions are currently being supported) can serve as a starting point for, or can be incorporated into, a dynamic representation. Dynamic representations can be fed by analytic monitoring; conversely, requirements for information produced by analytic monitoring can be driven by the need to maintain a current representation. Dynamic representations support situation awareness, and thus inform adaptive response and coordinated defense.

Dynamic representations can be used to enhance understanding, particularly of dependencies among cyber and non-cyber resources; validate the realism of courses of action; raise awareness of cyber threats, and support training and preparation; and identify gaps in planning, for which additional cyber courses of action need to be developed. Dynamic representations can include simulation exercises as well as executable models[40]. Models of adversary behavior can be game-theoretic.

Dynamic Representation involves coordination of, and fusion of information from mechanisms across architectural layers; therefore, no discussion of applicability to different layers is meaningful. See Appendix F for examples of specific mechanisms and their relative maturity. Cyber situation awareness is an area of active and extensive research, development, and transition activities; a survey of work in the cyber situation awareness area is beyond the scope of this report.

Dynamic Representation techniques rely on

- Information about systems and components that is also used by system, network, and security managers (e.g., configuration, security patch status, availability and performance statistics). Some information is provided by system or network management tools

[40] Simulation exercises can be model-based and automated, partially automated (e.g., training simulators, exercises, technology demonstrations), or purely manual (e.g., tabletop exercises). Simulation exercises are an established part of business continuity and disaster recovery [165]. Such activities can also lead to changes in organizational behavior, due to increased awareness [165].

(including available products for dynamic network mapping); other information may be provided by continuous monitoring or other security management tools.

- Information about functional dependencies among systems, networks, and components. This information is typically included in continuity or contingency planning documentation, where it becomes quickly outdated. Products are available for discovering IT assets and performing dependency mappings.[41]
- Information about mission dependencies on systems or services, networks or communications links, and information stores. This information may be provided by a Mission Impact Analysis or Business Impact Analysis (e.g., using the Map the Mission process [61], Mission Based Analysis [101], or Mission-Driven Assessment [102]). Modeling and automation remain research areas [103].

D.8 Non-Persistence

Non-Persistence techniques *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.[42] Virtualization makes such refreshing much easier. Non-Persistence is most appropriate when 1) refresh is quick enough not to interfere with mission/business functions, and 2) the elimination of services, information and connectivity is sufficient as to prevent an adversary from achieving their goals.

The effectiveness of Non-Persistence is enhanced by combining it with Diversity and Unpredictability. Substantiated Integrity (e.g., tamper-evident) mechanisms support ensuring that the information and services which are used to refresh have not been corrupted.

To maximize cost effectiveness, non-persistence (especially across multiple platforms) may be centrally managed. This central management may be viewed by an issue for some users/organizations where individual control of the platforms is the norm. Practical issues for managing server virtualization [104] need to be resolved in a manner consistent with QoS and security requirements.

D.8.1 Specific Techniques

D.8.1.1 Non-Persistent Information

When non-persistence is applied to information, the information is *refreshed* to a known trusted state and *deleted* when no longer needed. The deletion of information limits the opportunity of the adversary to exfiltrate critical information. The refreshing of information limits the opportunity of the adversary to modify critical system or mission information (resulting in corruption of services).

[41] Providers include BMC, HP, IBM, and VMware.

[42] Note for some situations both of these may not be required. For example, in a tactical environment it could be that even if the elimination of the information is not complete, the time required for an adversary to employ it in a meaningful way is too long for them to interfere with or otherwise adversely impact the mission.

The primary application of non-persistence for cyber resiliency is to information that is part of the state of a running process, whether stored in memory or written temporarily to storage media. Deletion (or elimination) of information written to storage media can involve multiple mechanisms. Automatic data deletion mechanisms in large-scale distributed processing [105] [106] and cloud environments [107] have been investigated, but are not yet available as products. Deletion of data from information stores remains problematic [108].

Media sanitization for magnetic storage media is very mature, but the same is not true for sanitization of solid state drives, especially for sanitization of an individual files [109]. Even the existing proven sanitization techniques may not provide a sufficiently rapid refresh capability. Some techniques may provide a sufficiently rapid refresh capability, but the data elimination may not be complete.

Table 31. Applicability and Maturity for Non-Persistent Information

Application Domain	Maturity
Hardware/physical storage media	Highly Mature for some hardware/physical storage media; immature for others.
Networking/communications	Relevant to cached information at network nodes; see System/network component.
System/network component	Transitional-to-Mature.
Operating system	Highly Mature.
Cloud, virtualization, and/or middleware infrastructure	Immature.
Mission / business function application / service	Dependent on application/service.
Information stores	Transitional in the context of DBMS. Immature for unstructured information stores, where responsibility devolves to the underlying operating system.
Information streams / feeds	See Non-Persistent Connectivity, below. Observation for Networks/communications above also applies.
System / system-of-systems	Immature.

D.8.1.2 Non-Persistent Services

Non-Persistence applied to services defends against malware insertion, and increases the adversary's work factor for establishing a foothold. Running an end-user system off bootable media provides non-persistence. When non-persistence is applied to services (on a client or end-user device, on a server), services are refreshed and are terminated when no longer needed. The refreshing of the services limits the window of opportunity for the adversary to implant malware. Stateful refreshing of services can be combined with diversity [33]. Automatic termination of applications on end-user devices is part of some operating systems.

Table 32. Applicability and Maturity for Non-Persistent Services

Application Domain	Maturity
System/network component	Mature; Transitional in the cyber resilience context.
Operating system	Highly Mature, including non-persistent virtual desktops.
Cloud, virtualization, and/or middleware infrastructure	Immature.
Mission / business function application / service	Dependent on application/service.

D.8.1.3 Non-Persistent Connectivity

Non-Persistence applied to connectivity defends against sniffing, eavesdropping, and intelligence-gathering (e.g., network mapping). When non-persistence is applied to connectivity (at the application or network layer), connections are refreshed and are terminated when no longer needed. Applications typically can tolerate dropped connections to services on which they rely, and seek to reinstate those connections. Some applications (particularly those designed for mobile devices) terminate connections they no longer need, to improve performance. Thus, applications could generally be expected to tolerate refreshing or terminating connections to change the attack surface.

At the network layer, ports and protocols can be terminated when an organization determines that they are unneeded; however, this is a relatively static configuration change. Products for managing dynamic connections (e.g., via DHCP) are relatively mature. However, the goal of such products is to improve performance, rather than to terminate or refresh connections to change the attack surface.

Finally, specific information-gathering devices (e.g., sniffers, packet capture devices) should be connected to the network only as long as they are being used; otherwise, they become high-value targets.

D.9 Privilege Restriction

Privilege Restriction techniques *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*[43] *and trust*[44] *respectively, to minimize the potential consequences of adversary activities*. Generally, the more critical the asset the more fine-grained the privileges that may be applied to it, and the more trusted an entity is, the greater privilege that it is granted. Privilege Restriction must be aligned with Coordinated Defense, to ensure that privileges are defined and managed consistently across architectural layers and across systems in a system-of-systems.

D.9.1 Existing and Emerging Technologies

Examples of privilege restriction mechanisms include fine-grained access control and trust-based privilege assignment. Access control mechanisms are a mature technology, as are privilege mechanisms. Effective management of privileges remains challenging, due to the size, complexity, and dynamic nature of many enterprises. Products and toolsets are available to visualize the relationships among roles, responsibilities, and privileges, and to define and enforce least privilege. These include tools based on frameworks for federated identity and privilege management. Thin clients enable privilege restrictions to be applied on servers rather than on end-user devices. However, many thin clients depend on vulnerable browsers to deliver functionality to end users.

Dynamic or risk-adaptable privilege management remains a research area, with risk-adaptable access control (RAdAC) mechanisms becoming transitional [110].

[43] Criticality is determined based on analysis of the potential consequences of compromise; higher criticality requires more restrictive (typically more fine-grained as well as more closely controlled) privileges.
[44] Trust in a user is determined based on organizational policies and practices; trust in a cyber-resource depends not only on which user (if any) it represents, but also on such factors as its provenance and its recent history.

D.9.2 Applicability and Maturity

The following table presents an overall assessment of the maturity of Privilege Restriction techniques and technologies, as applied to different architectural layers.

Table 33. Applicability and Maturity for Privilege Restriction Techniques

Application Domain	Privilege-Based Restrictions	Privilege Management
Hardware/firmware	Mature (multiple CPU modes or rings).	Mature (OS- or hypervisor-based).
Networking/ communications	Mature for address-based restrictions; Transitional for identity-based restrictions.	Mature for address-based restrictions; Transitional for identity-based restrictions.
Operating system	Mature.	Mature; Immature-to-Transitional for dynamic escalation of privilege restrictions.
Cloud, virtualization, and/or middleware infrastructure	Mature.	Mature; Immature-to-Transitional for dynamic escalation of privilege restrictions.
Mission / business function application / service	Dependent on application/service.	Dependent on application/service.
Information stores	Mature for DBMSs and shared data repository products.	Mature for DBMSs and shared data repository products.
System / system-of-systems	Mature, but requires consistent management.	Immature-to-Transitional, based on federated identity and privilege management frameworks.

D.10 Realignment

Realignment techniques *align cyber resources with core aspects of mission/business functions, thus reducing the attack surface associated with resources dedicated to less significant activities.* Realignment entails defining, and determining the operational implications and cyber resource needs of, alternative as well as primary mission and cyber defender courses of action. 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 off-loading some less important cyber-supported functions to a service provider that is better able to support the functions.[45] Other examples are to perform a function using out-of-band communications (e.g., replace automated cross domain services with air gaps and sneaker nets), or to eliminate 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. Realignment can also involve re-implementation or custom development of critical components.

Realignment involves reallocation of resources across architectural layers; therefore, no discussion of applicability to different layers is meaningful.

D.11 Redundancy

Redundancy techniques *maintain multiple protected instances of critical resources (information, connectivity, 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.

[45] The trust required of the provider depends upon the importance of the functions and the sensitivity of the data it must handle.

Maintaining an instance means keeping it compliant with the requirements that apply to it (e.g., patching software for security, updating databases for data quality), whether or not it is actively used.[46] Maintaining a *protected* instance of a critical resource means viewing each instance as an adversary target and recognizing and mitigating ways in which a successful attack on one instance could propagate to all instances.

Redundancy is a commonly cited technique for resilience in general [67]; however, some attention has been paid to the cyber threat [68]. Diversity is vital to the effectiveness of Redundancy (e.g., instances can provide the same functionality or information, while being implemented in different ways). Redundancy can be most effective in combination with Privilege Restriction (to ensure that all copies are protected consistently), Analytic Monitoring (to ensure that all copies are monitored consistently), and Segmentation (e.g., instances can be protected by placing them on different segments).

D.11.1 Existing and Emerging Techniques

Redundancy is a mature technique at multiple layers, for example:

- Redundant hardware, redundant copies of software and data, backup and restore procedures, and failover to backup facilities are well-understood aspects of contingency planning. In addition, redundancy can be applied at the chip level [111].
- For networking, redundancy involves providing alternative communications paths. Network topologies typically provide multiple paths. Products and services for fault-tolerant networking, particularly for industrial control systems (ICS), are available [112], and can be integrated with diversity to make use of different communications media [113].
- Redundant data feeds are incorporated into some real-time architectures, e.g., for financial systems [114].

In cloud environments, Redundancy can be used to improve fault-tolerance for data stores [115], and can be combined with Non-Persistence [34].

Determination of whether multiple instances of the same resource are truly redundant presents challenges. For networking and for distributed processing, insight into whether a single point of failure exists can be limited, either because networking or processing services are provided by an external service provider [116], or due to limitations in mapping and analysis tools. For software, the challenge is to ensure that redundant copies have the same patch status and configuration; latency in patching can be difficult to ascertain. Similarly, for information stores, latency in updates can result in inconsistencies among apparently duplicate stores.

D.11.2 Applicability and Maturity

The following table presents an overall assessment of the maturity of Redundancy techniques and technologies, as applied to different architectural layers.

[46] Redundancy can be implemented in multiple ways: active redundancy, in which redundant components are fully operational; standby, in which redundant components are partially activated; and passive, in which redundant components are off-line. Maintaining protected instances is more challenging for passive redundancy.

Table 34. Applicability and Maturity of Redundancy Techniques

Application Domain	Specific Technique(s)	Assessment
Hardware/firmware	Duplicate hardware Redundancy at the chip level	Mature, but costly; need to consider SCRM Transitional
Networking/communications	Alternate communications paths	Mature for network communications (particularly for ICS)
System/network component	Duplicate hardware/software	Mature, but costly; maintaining consistent and current patch status & configuration presents challenges
Mobile system/network component	Duplicate hardware/software Wireless backup	Mature, but costly; maintaining consistent and current patch status & configuration presents challenges Transitional-to-Mature for wireless backup; bandwidth and connectivity can be problematic.
Operating system	Duplicate copies of installable image	Mature; maintaining consistent and current patch status & configuration presents challenges
Cloud, virtualization, and/or middleware infrastructure	Extra capacity	Mature; insight into degree of redundancy may be limited
Mission / business function application / service	Redundant copies of running software Backup & restore capabilities	Mature, but can impose operational costs Mature, part of standards of good practice
Information stores	Database replication Backup & restore capabilities	Mature, part of standards of good practice
Information streams / feeds	Redundant data feeds	Mature for some sectors (e.g., financial)
Systems	Backup & restore capabilities	Mature, part of standards of good practice

D.12 Segmentation

Segmentation and isolation techniques *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. Segmentation can enable resources to be isolated via dynamic reconfiguration, as part of adaptive response.

D.12.1 Existing and Emerging Technologies

Segmentation often employs either physically distinct entities or virtualization of computing enclaves to provide the desired separation. Segmentation can be applied within a component, with hardware support [117]. Encryption can be used to define different segments within a network.

Defining enclaves or sub-networks within an intranet is an established practice. Of particular interest for cyber resiliency is placing an organization's Security Operations Center (SOC) on a separate sub-network. However, such placement could restrict what tools within the SOC can observe from other sub-networks. In addition, physical separation of sub-networks is challenging to achieve in the current technology environment, as devices are increasingly enabled for wireless communication and as reliance on common communications infrastructures increases.

Other established practices include isolating an intranet from an extranet, and both from the Internet, separating inbound from outbound traffic, and separating requests from responses. Segmentation can also be applied at the system layer, by using virtualization, at the application

layer, by partitioning services, and at the data layer, by providing separate data repositories (e.g., based on provenance, in conjunction with Substantiated Integrity).

D.12.2 Applicability and Maturity

The following table presents an overall assessment of the maturity of Segmentation techniques and technologies, as applied to different architectural layers.

Table 35. Applicability and Maturity for Segmentation Mechanisms

Application Domain	Maturity
Hardware/firmware	Highly Mature for some hardware using ring mechanisms (see Privilege Restriction), but not dynamic.
Networking/communications	Highly Mature, using firewalls or cross domain solutions (CDS) to define enclaves. Immature-to-Transitional for placement of Security Operations in a separate enclave.
System/network component	Highly Mature for some components, using Privilege Restriction or physical separation of processing planes; Immature-to-transitional for others.
Operating system	Highly Mature.
Cloud, virtualization, and/or middleware infrastructure	Immature-to-Transitional. While a virtualization infrastructure manages separate virtual machines, attacks can circumvent those mechanisms.
Mission / business function application / service	Immature; typically relies on underlying operating system.
Information stores	Mature using encryption. Immature-to-Transitional for pedigree-based segmentation.
Information streams / feeds	Highly Mature, using cryptographic separation.
System / system-of-systems	Transitional, using multi-level security (MLS) or multiple security levels (MSL) approaches.

D.13 Substantiated Integrity

Substantiated Integrity techniques *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., cryptographic seals or checksums on critical records or software), data validation (checking that data conforms to its specified requirements, such as type or range), program verification [118], polling of inputs from diverse critical services (e.g., Byzantine quorum systems) to determine correct results in case of conflicts between the services, and tamper-evident technologies.

D.13.1 Existing Techniques and Technologies

A variety of existing techniques are highly mature, for example:

- Software and Data Integrity Checks: A variety of tools, many integrated into operating systems, enable checksums – particularly cryptographic checksums or seals – to be applied to software, critical files and records, and data in transit. Multiple COTS tools can provide cryptographic checksum integrity checking to backups as well as primary versions [119] [120].[47] Integrity checks for virtual/cloud environments are immature-to-transitional [121], as they are for software on mobile devices . Trusted path mechanisms provide confidence that information exchanged by an end user and an operating system

[47] Digital fingerprinting provides a mechanism for identifying and tracking content [165].

has not been corrupted, usually with hardware support. Software-mediated trusted path mechanisms at the application layer are transitional [122], as are mechanisms that operate in virtual environments [123].

- Network Address Validation: Multiple products provide mechanisms for validating packets, to provide dynamic ARP (Address Resolution Protection) protection.
- Data Validation: DBMSs enable constraints to be applied to ensure that the data remains consistent with quality criteria. Integrity checking techniques can be extended to increase the correctness and efficiency of established constraint-checking [124]. Commercial and open source [125] tools are available for data validation for Web applications.

At the hardware level, some tamper-evident technologies are mature [126], while others are transitional. Substantiated Integrity for firmware is transitional [127] to mature [128].

D.13.2 Emerging Techniques and Technologies

D.13.2.1 Data Provenance and Trust

The W3C Provenance Incubator Group [129] has offered the following definition of provenance:

> Provenance of a resource is a record that describes entities and processes involved in producing and delivering or otherwise influencing that resource. Provenance provides a critical foundation for assessing authenticity, enabling trust, and allowing reproducibility. Provenance assertions are a form of contextual metadata and can themselves become important records with their own provenance.

The W3C Provenance Incubator Group report provides extensive analysis of the very active provenance research area, and multiple recommendations. Mechanisms are being explored to provide capabilities to establish data provenance and trustworthiness [130] [131] [132]. Research into provenance analytics has been recommended [133], and the relationship between data provenance and trust is a challenge area [134]. Provenance as a basis for access control is being explored [135], and requires protection of provenance assertions.

D.13.2.2 Byzantine Quorum Systems

Byzantine fault tolerance has been extensively researched, but most work remains transitional. The application of Byzantine fault tolerance to intrusion tolerance remains immature, but of high interest [136] [137]. Application to cloud computing is in initial stages [138], but can leverage multiple cloud systems [139].

D.13.3 Applicability and Maturity

The following table presents an overall assessment of the maturity of Substantiated Integrity techniques and technologies, as applied to different architectural layers.

Table 36. Applicability and Maturity for Substantiated Integrity Mechanisms

Application Domain	Maturity
Hardware/physical storage media	Highly Mature for some hardware/physical storage media, using tamper-evident technology; immature for others.
Networking/communications	Highly Mature when encryption is used; Mature for some mechanisms (e.g., ARP protection); Transitional in the context of IPv6; otherwise Immature.
System/network component	Highly Mature for some components, using tamper-evident technology; Immature-to-Transitional for firmware; Immature for many components.
Operating system	Highly Mature.
Cloud, virtualization, and/or middleware infrastructure	Immature-to-Transitional.
Mission / business function application / service	Immature; typically relies on underlying operating system.
Information stores	Immature-to-Transitional in the context of DBMS. Immature for unstructured information stores.
Information streams / feeds	Highly Mature when encryption is used; Highly Mature for some trusted path mechanisms, Transitional for software-based trusted path.
System / system-of-systems	Immature, particularly in the cyber resilience context.

D.14 Unpredictability

Unpredictability techniques *make changes frequently and randomly*, *not just in response to actions by the adversary*. These changes, which may draw upon Diversity, Non-Persistence, and Dynamic Positioning practices, make it more difficult for an adversary to predict behavior and (as with Coordinated Defense) this increases the chance of adversary actions being detected or tradecraft revealed. Examples of unpredictable behavior include, but are not limited to, address space layout randomization (ASLR), changing browsers and authentication mechanisms, encryption rekeying, and changing permitted ports. Unpredictability is intended to be transparent to end users. A key challenge is providing enough unpredictability to achieve the intended benefits [140].

Appendix E POET Considerations

Adoption and effective use of a cyber resiliency technique entails addressing a variety of challenges. For each technique, factors affecting adoption and use are identified using the POET (political, operational, economic, and technical) approach [141] [133]. POET factors largely operate to limit or shape adoption or use, but can also serve to increase uptake. The general identification of POET considerations (i.e., how the general POET factors apply in the context of a given technique) in this appendix can serve as a starting point for identification and analysis of more specific factors as part of developing recommendations.

Table 37. General POET Factors

General POET Factors	
Political: • Policies, laws, and regulations, which may constrain the use of some techniques or solutions • Relationships and commitments, including service level agreements • Governance, including which roles, responsibilities, and processes have been defined • Risks and risk tolerance, including reputation risk and tolerance for mission and programmatic risks • Organizational culture • Investment strategy, which may include a commitment to technical standards or product suites, and may affect time-phased application of solutions	**Operational:** • Mission priorities • Mission impacts, particularly as mission dependencies on cyber resources change • Operational constraints, including physical constraints (footprint, power, connectivity) in operational environments • Impacts on supporting processes, including security management, cyber defense, and provisioning • Flexibility/agility, including ability to adapt to changing mission needs
Economic: • Costs, including life-cycle and staffing costs • Benefits, including reduced costs or increased opportunities • Perceived value • Incentives	**Technical**: • Standards, including restrictions imposed by existing standards as well as the absence of agreed-on standards • Performance • Legacy investments • Interoperability • Infrastructure

Table 38. POET Considerations for Adaptive Response

Adaptive Response	
Political: • Lack of clarity regarding governance or decision-making authority, particularly when responses span management/ownership domains • Concern for adherence to SLAs • Concern for liability in case of collateral damage	**Operational:** • Integration of processes and procedures to apply existing reconfiguration and resource allocation mechanisms into CCoAs as well as SOPs • Potential instability arising from poorly coordinated human response activities at different architectural layers or for different systems in a system-of-systems • Potential instability arising from purely automated dynamic response mechanisms • Potential lack of visibility into purely automated response mechanisms
Economic: • Costs of acquiring, integrating, and maintaining the mechanisms as part of systems and components • Benefits of more reliable service • As dynamic mechanisms become more robust, decreased need for operator oversight and intervention	**Technical**: • Interoperability of Adaptive Response mechanisms across architectural layers and across systems • Differences in maturity of solutions

Table 39. POET Considerations for Analytic Monitoring

Analytic Monitoring	
Political: • Potential for abuse of large volumes of monitoring data, particularly when aggregated (e.g., privacy, sensitivity due to determination of organizational or mission priorities and plans based on inference and aggregation) • Concerns for reputation and liability, when analysis involves sharing information with other organizations • Integration of cyber and non-cyber monitoring, since they are often under different reporting chains • Trade-offs between insights gained from monitoring and protection from encryption	**Operational:** • Coordinating monitoring across architectural layers and across systems • Sharing information and coordinating analysis, particularly in support of damage assessment • Constraints associated with operational environment (e.g., bandwidth, power, or storage limitations) • Lack of visibility into non-owned infrastructures (e.g., networks, cloud computing environments)
Economic: • Costs of maintaining and protecting monitoring data • Costs of expert analysts, particularly for malware and forensic analysis	**Technical**: • Data interoperability, to fuse and analyze monitoring data across architectural layers and across systems • Monitoring in virtualized or constrained environments • Need for analytic capabilities, whether in the form of malware analysis, red teaming, or damage assessment, to keep pace with changes in adversary capabilities as well as in enterprise information and communications technologies

Table 40. POET Considerations for Coordinated Defense

Coordinated Defense	
Political: • Governance (e.g., who identifies roles and responsibilities; oversight to ensure that activities are coordinated; which responsibilities are retained and which may be outsourced, for example to a CND Service Provider) o Coordination of security management, network management, and system management activities often not part of staff job descriptions o Liability (e.g., who will be held accountable when something goes wrong) • Information sharing to support coordination, particularly when the shared information reveals weaknesses or gaps in an organization's or a business unit's governance	**Operational:** • Presentation of information (visualization, ensuring that staff are presented with information at a level appropriate to their responsibilities) • Lack of visibility into effects of actions on non-owned resources (e.g., hosted or managed services) • Staffing o Overloaded roles (e.g., administrator with cyber incident responsibilities as additional duties) o Difficulties obtaining and retaining staff with the needed expertise o Lack of training that includes potential unintended consequences of administrator or defender actions • Knowledge capture and development of CCoAs. • Making cyber defense part of mission exercises, as well as exercises of contingency plans
Economic: • Costs of providing resources (cost of multiple defense measures, cost of ensuring the multiple measures are coordinated), particularly staffing	**Technical**: • Automation and visualization, so that decision makers at all tiers are presented with actionable information • Automated means of detecting conflicts between measures

Table 41. POET Considerations for Deception

Deception	
Political: • Reputation and relationship risks. To effectively fool the APT, deception environments must be realistic. To achieve that may involve deceiving non-hostile users, or even the public, which can raise public relations(PR) concerns. • Potential legal, regulatory, contractual, or policy constraints on employing deception measures • OPSEC risks. A realistic deception environment may enable an adversary to infer sensitive information. • Liability risks	**Operational:** • Challenge of maintaining realism of deception environments (e.g., honeynets, deception databases) and deceptive information (e.g., fabricated system logs, deceptive information on adversary-accessible sites) • Operational integration of deception into the organization's defense • Operational challenges to monitoring created by encryption of network traffic
Economic: • Costs of maintaining viable deception environments and deceptive information • Cost of developing realistic deception scenarios and staff to respond in realistic manner (currently time/labor intensive)	**Technical**: • Challenges to creating and maintaining a viable deception net that can deceive an adversary for a considerable period of time, particularly on a realistic scale

Table 42. POET Considerations for Diversity

Diversity	
Political: • Organizational policies requiring adherence to an enterprise architecture (possibly including restriction to a specific set of software products) • Trend toward Bring Your Own Device (BYOD), which increases incidental diversity	**Operational:** • Maintaining an accurate representation of enterprise systems • Consistent management and use (particularly when different instances of the same capability present different user and/or administrator interfaces) • Maintaining IT and help desk support for diverse set of services
Economic: • Increased acquisition/procurement costs related to SCRM • Increased life-cycle costs to acquire, operate, and maintain multiple instances of the comparable capability • Cost to discover legacy applications that need to be retooled so as to work across multiple instances of capabilities • Increased cost of training personnel on multiple instances of capabilities	**Technical**: • Ability to determine whether two designs or implementations are truly different • Interoperability o Specifications and standards for interfaces between architectural layers and for functional capabilities o Validation: Even with clear specifications, different versions can fail to interoperate correctly, so test cases (particularly cases that can be applied relatively early in the development and integration process) are vital

Table 43. POET Considerations for Dynamic Positioning

Dynamic Positioning	
Political: • Concerns for meeting service level agreements when dynamic repositioning is used	**Operational:** • Potential lack of visibility into location of resources that, due to mission needs, are currently mission-critical and require heightened protection • Potential limitations due to operational environment (e.g., need to retain processing within an enclave)
Economic: • Cost to migrate mission applications to distributed processing environments • Costs of distributed processing and distributed data • Potential gains in efficiency, particularly for Analytic Monitoring when sensors can be dynamically positioned	**Technical**: • Maintaining consistency and integrity for distributed processing and distributed data • Potential performance impacts • Technical limitations due to policy or programmatic restrictions (e.g., organizational commitment to a specific product or product suite which does not accommodate repositioning) • Immature-to-transitional technologies

Table 44. POET Considerations for Dynamic Representation

Dynamic Representation	
Political: • Sensitivity of derived information, particularly about mission dependencies and current adversary characteristics / behaviors • Governance issues / information sharing constraints in the context of systems-of-systems	**Operational:** • Lack of visibility into non-owned infrastructures (e.g., networks, cloud computing environments) • CONOPS/SOPs/rules of engagement to configure and tune mapping tools appropriately, and grant cyber defenders access to their results
Economic: • Increased costs to acquire, use, and maintain tools • Improved efficiency in resource management	**Technical**: • Moving to near-real-time representations • Data interoperability for correlating and fusing information from multiple tools (including those used for Analytic Monitoring)

Table 45. POET Considerations for Non-Persistence

Non-Persistence	
Political: • Forensics. Some organizations need to perform digital forensics to help identify the nature of adversary malware. As an image is refreshed, legal evidence that might be needed for prosecution could be eliminated. So refresh in many instances will have to be accompanied by ensuring that sufficient information (snapshots or images) is captured and retained that is sufficient for legal purposes.	**Operational:** • Management. To maximize cost effectiveness, non-persistence (especially across multiple platforms) may need to be centrally managed.
Economic: • Initial costs of establishing non-persistence • Cost of retooling existing services and applications to support non-persistence	**Technical**: • Quality of service (QoS). For some mission/business functions, the refresh capability needs to be relatively seamless to ensure that it does not disrupt, or minimizes the disruption of, organizational operations. • Speed. Deletion and sanitization technology of storage media is generally not rapid enough or applied across broad enough spectrums of media to provide effective non-persistence. • Virtualization. No specific technical barriers for employing virtualization technology. But virtualization in support of non-persistence will need to be applied to a broader range of products and devices, for some of which there are no existing non-persistence activities to build upon.

Table 46. POET Considerations for Privilege Restriction

Privilege Restriction	
Political: • Governance and CONOPS issues (e.g., inconsistencies or gaps in definitions of roles, responsibilities, and related privileges) • Asset criticality determination may be impacted by political concerns of competing organizational entities. • Organizational cultures in which staff expect to have full control over the cyber resources they use.	**Operational:** • Identifying resource criticality (and possibly reflecting changes to resource criticality as mission priorities change) • Managing privileges in changing mission contexts • Operational impetus to share roles
Economic: • Management costs • Staffing costs (e.g., cost for imposing dual-authorization on what were single roles)	**Technical**: • Absence of standards and reference implementations for dynamic (risk-adaptable) privilege restriction • Absence of generally-accepted standards and reference implementations for federated privilege management

Table 47. POET Considerations for Realignment

Realignment	
Political: • Organizational and cultural impacts (e.g., eliminating functions that personnel are used to employing, impact on morale of relocating staff)	**Operational:** • Lack of visibility into non-owned infrastructures can reduce the effectiveness of other techniques • Operational convenience • Operational disruption due to changes in processes and procedures • Potential for a performance hit in those instances where automated process/communications are replaced by more secure, but less rapid procedural measures (e.g., replacing direct communications with sneaker net)
Economic: • Cost of performing the analysis of determining functions that need to be realigned • Resource reallocation to pay an external entity to provide for services that previously were covered as part of organizational costs (can increase or decrease costs, but either way resources are reallocated)	**Technical**: • Limited existing automated means to determine which existing functions are a potential attack vector of an adversary • Co-mingling of IT and software that support core mission functions and secondary functions may make it more challenging to remove the secondary function

Table 48. POET Considerations for Redundancy

Redundancy	
Political: • Organizational commitment (for cost and other reasons) to a limited set of service providers (e.g., for communications, backup, or cloud services)	**Operational:** • Limited visibility into the resources provided by service providers, which makes validation of redundancy problematic
Economic: • Costs of acquiring, maintaining, and securing redundant resources	**Technical**: • Timing, to ensure that redundant capabilities are either equally up-to-date, or that differences in how current different copies are known

Table 49. POET Considerations for Segmentation

Segmentation	
Political: • Trend toward a seamless computing and network infrastructure • Legacy policy of defining enclaves based solely on the basis of confidentiality	**Operational:** • Modification of CONOPS, administrator SOPs, and cyber CCoAs • Limited visibility across segments o Placing an organization's SOC on a separate sub-network can limit visibility
Economic: • Cost and schedule impacts of re-architecting o Costs associated with physically separating key components / networks (e.g., additional firewalls, routers) o Costs associated and with applying virtualization to segmentation	**Technical**: • Dynamic ability to segment different parts of an architecture supported by virtualization. Experimentation is needed to demonstrate and validate the concept.

Table 50. POET Considerations for Substantiated Integrity

Substantiated Integrity	
Political: • Policy support for SCRM/AT	**Operational:** • User confidence and ability to act (requires development of alternative mission courses of action based on awareness that some cyber resources cannot be trusted)
Economic: • Cost and schedule impacts (e.g., of incorporating and managing cryptographic checksums on data)	**Technical**: • Use of polling between distributed services (e.g., byzantine quorum) is generally not available in COTS products • Limited availability of Trusted Platform Module (TPM)-enabled products

Table 51. POET Considerations for Unpredictability

Unpredictability	
Political: • Frequent and unanticipated changes may require approvals from multiple authorities • Impact on organizational culture, if unpredictability is not user-transparent	**Operational:** • Operational/mission impacts. Unplanned changes can adversely impact ongoing operations that assume predictability and stability of key activities. • Administrator/help desk impacts. If the unpredictability is not transparent to end users, help desk load will increase.
Economic: • Some changes associated with unpredictability are still done manually, thus more frequent and unanticipated changes may add additional cost to operations.	**Technical**: • Technical challenge of validating that unpredictable behavior is user-transparent • Unpredictability is closely linked to non-persistence, dynamic positioning, and diversity and therefore shares the technical limitations of all of those techniques.

Appendix F Time-Phasing of Cyber Resiliency Solutions

The following table presents an initial set of representative examples of cyber resiliency practices in the near-, mid-, and long-term. The time periods are determined by technical maturity, i.e., by how soon the specific technology or technology-dependent procedures could be integrated into an operational architecture. The set of examples is expected to change over time. The criteria for assigning an example to a time period are:

- Near-term (within than 3 years): Operating procedures and/or commercial solutions are available today, or will be readily available for integration into an architecture within two years (and hence can be integrated within 3 years).
- Mid-term (within 5 years): Procedures and/or commercial solutions are expected to be available within 3-4 years, or exist but present scaling or other practical constraints that limit integration and deployment in the near term.
- Long-term (more than 5 years): The technology or technology-dependent procedure requires additional R&D to determine its viability, or has been demonstrated only on such a scale or in such a constrained environment that it cannot be integrated and deployed in the mid-term.[48]

[48] For a survey of the cyber resiliency research landscape, see [10].

Table 52. Representative Examples of Cyber Resiliency Mechanisms

Technique	Near-Term	Mid-Term	Long-Term
Adaptive Response: Take actions in response to indications that an attack is underway based on attack characteristics	• Manually administered reconfiguration (e.g., restriction of ports or protocols, termination of selected services) • Manually administered re-provisioning / reallocation of resources (typically using virtualization) • Development of policies to support automated dynamic reconfiguration when technologies become available	• Development and testing of applications/services that have been previous presented to adversaries (deploy when conditions warrant) • Incorporation of running desktops/laptops from bootable media as part of standard attack response	• Automated dynamic reconfiguration (e.g., restriction of ports or protocols, termination of selected services) • Automated dynamic re-provisioning / reallocation of resources (typically using virtualization) • Dynamic reconstitution
Analytic Monitoring: Gather and analyze data on an ongoing basis and in a coordinated way to identify potential vulnerabilities, adversary activities, and damage	• Perform malware and forensic analysis • Penetration testing / active probing • Cyber sensor data fusion (within system, across systems) • Mechanisms to make encrypted traffic visible to monitoring (selected protocols)	• Monitoring and probing of supply chain defenses • Limited integration of monitoring in virtualized / cloud environments with security monitoring • Coordinate sensor coverage to avoid gaps or blind spots • Analyze data to identify anomalies, develop I&W, and monitor effectiveness of cyber courses of action • Mechanisms to make encrypted traffic visible to monitoring (additional protocols)	• Near-real time forensic analysis • Automated damage assessment based on data quality / pedigree mechanisms • Security monitoring in virtualized / cloud environments • Dynamic management of cyber sensors (e.g., increased sensing at key choke points, reposition or reconfiguration of sensors)

Technique	Near-Term	Mid-Term	Long-Term
Coordinated Defense: Manage adaptively and in a coordinated way multiple, distinct mechanisms to protect critical resources from adversary activities	• Coordinated placement of cyber sensors • Use of a defense in depth strategy within organization • Provide equivalent security capabilities from different vendors in different locations (e.g., anti-virus on organization firewall, servers, and desktops)	• Employ a systematic process to identify dependencies and interactions among cyber defenses, security controls, and performance controls • Coordinated definition of privileges • Response activities coordinated to avoid mission impairment, blinding of cyber defenders • Recovery activities coordinated to avoid allowing gaps in security coverage	• Automated identification of conflicts and dependencies among defenses • Integration with diversity and non-persistence to provide defense in depth with diverse and dynamically changing methods • Automated support for cross-organizational coordination
Deception: Use obfuscation and misdirection (e.g., disinformation) to confuse an adversary	• Encryption of mission data • Honeypots (low interaction, based on commonly used attacker requested services) • Encryption of security control information	• Use of routers and firewalls to hide sensitive subnets[49] • Seeding sites known to be frequented by adversaries with false information regarding an organization's security posture[50] • Honeynets (network of honeypots intended to imitate activities of a real system)[51]	• Use of honeynets and virtualization to run deception nets that respond dynamically to actions of an adversary and can do so for an extended period of time • Use of fabricated system logs and/or security management documentation

[49] This is technically feasible even in the near-term, but standards of practices have not been established, and would need to be integrated with architectural and operational strategies for managing, configuring, and maintaining the security posture of hidden cyber resources.
[50] This is technically feasible even in the near-term, but is operationally challenging to carry out in a credible way without inadvertently disclosing real information. Standard or generally accepted operating procedures are not expected to be available within 3 years.
[51] Honeynets already exist. But to be effective they need to be incorporated into an organization's cyber security operations center and integrated with analytic monitoring capabilities.

Technique	Near-Term	Mid-Term	Long-Term
Diversity: Use a heterogeneous set of technologies (e.g., hardware, software, firmware, protocols) and data sources to minimize the impact of attacks and force adversaries to attack multiple different types of technologies	• Different browsers on operating systems (OSs) • Limited diversity of operating systems • Diversity of apps on smartphones and tablets	• Use of different protocols / communications diversity (e.g., over time, space, frequency) • Diverse suite of platforms for end users (e.g., some using tablets, some laptops) • Diverse mechanisms for critical security services, e.g., authentication • Use of different suppliers of critical components in supply chain	• Hardware diversity via custom chip sets • Determinable degree of data diversity (e.g., pedigree-based) • Dynamically employ different OSs and different applications on laptops, desktops and servers (virtualization-enabled linkage of non-persistence and diversity) • N-version programming of applications • Use of obfuscating and randomizing compilers • Tailored compiling of applications and OSs
Dynamic Positioning: Use distributed processing and dynamic relocation of critical assets and sensors	• Changing IP and MAC addresses	• Periodic repositioning of key organizational data	• Dynamic repositioning or provisioning of cyber sensors • Moving target defense (state capture, state restoration on a different platform), frequently enabled by platform / OS independence (e.g., TALENT)
Dynamic Representation: Construct dynamic representations of components, systems, services, adversary activities, and effects of alternative cyber courses of action	• Dynamic network mapping tools • Dynamic asset discovery and dependency mapping tools • Reality-based security awareness training (e.g., recognizing and reporting potential indicators of insider threat)	• Correlation of information from monitoring tools employed throughout the information system to achieve organization-wide situational awareness	• Near-real-time adversary modeling • Near-real-time modeling of functional dependencies and security / performance posture • Near-real-time modeling of mission dependencies on cyber resources

Technique	Near-Term	Mid-Term	Long-Term
Non-Persistence: Retain information, services, and connectivity for a limited time, thereby reducing an adversary's opportunity to exploit vulnerabilities and establish a persistent foothold	• Desk top virtualization • Running desktops/laptops from bootable media • Server virtualization (e.g., email and DNS servers) • Data non-persistence with assured deletion for widely used platforms / storage media (e.g., magnetic media sanitization mechanisms) • Automated termination of – or state capture and restart for – services • Policies, procedures, and mechanisms for terminating unused ports and protocols • Policies and procedures for connecting, using, and disconnecting sniffers and packet capture devices	• Integration with monitoring to refresh as conditions warrant • Applying virtualization to stateful services (e.g., active directory, routers) • Data non-persistence with assured deletion for newer storage media (e.g., solid state media sanitization mechanisms) • Data non-persistence with assured deletion / transformation for key platforms (e.g., automatic encryption of archives and logs for servers, routers) • Configuration of dynamic network management to terminate and refresh connections unpredictably or based on time or usage limits	• Non-persistence (media/device sanitization or data transformation via encryption) for smartphones and tablets • Integration of non-persistence (for information) and substantiated integrity • Coordinated application of non-persistence across a mission or enterprise to support the overall mission or enterprise
Privilege Restriction: 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	• Use of thin clients • Removal of admin rights from end users for their machines • Dual authorization for critical functions	• Separate processing domains based on privilege • Criticality-based restriction of privileges required to use resources • Trust-based assignment of privileges to users and cyber entities	• Dynamic escalation of privilege restrictions based on indications of adversary activities
Realignment: Align cyber resources with core aspects of mission/business functions, thus reducing the attack surface	• Identification and offloading of non-mission-essential functions • Clear definition of managed interfaces, with exfiltration protections	• Re-implementation or custom development of key critical software components • Consideration of attack surface reduction as part of systems engineering / development	• Re-implementation or custom development of key critical hardware components

Technique	Near-Term	Mid-Term	Long-Term
Redundancy: Maintain multiple protected instances of critical resources (information and services)	• Data backup • Hot/warm/cold backup services • Redundancy in supply chains	• Wireless backup of tablets • Alternate communications • Mission data stored in multiple locations	• Redundancy coupled with diversity for supporting infrastructures (e.g., multiple power suppliers)
Segmentation: Separate (logically or physically) components of dubious pedigree from more trusted ones, to limit the spread of or damage from successful exploits	• Defining enclaves or sub-networks within an intranet • Use of virtualization to separate application environments on different virtual machines • Use of routers/firewalls to separate Internet from Intranet • Use of routers to separate different parts of an organization's DMZ • Use of routers and other means to separate organization's security operations center from rest of network (isolate from attack) • Separate inbound traffic from outbound traffic	• Separate critical from non-critical data (e.g., using encryption) • Separate critical from non-critical processing via subnets supported by routers and firewalls	• Physically separate critical and non-critical services • Use of virtualization and/or encryption to separate critical and non-critical services • Dynamic isolation of sub-networks or sets of resources during an attack
Substantiated Integrity: Ascertain that critical services, information stores, information streams, and components have not been corrupted by an adversary	• Encrypted checksums on critical data • Use of WORM drives • Two-person control on critical changes • Tamper-evident / anti-tamper controls	• Employ and validate checksums on all retrievals of data in database • Validation of provenance of information provided by external entity • Trusted Platform Module (TPM)-based attestation and verification of boot processes	• Use of polling between distributed services (e.g., Byzantine quorum) to ascertain and implement "correct" action if an individual service is compromised by adversary
Unpredictability: Make changes frequently and randomly, not just in response to actions by the adversary	• Random change of crypto keys • Randomize communications patterns (e.g., frequency-hopping) • Randomization of session IDs • Centralized, automated frequent randomization of admin passwords	• Integration of unpredictability with non-persistence (e.g., refresh services at random intervals)	• Integration of unpredictability with diversity, randomly changing protocol suite • Integration of unpredictability with dynamic positioning, randomly moving services / applications

Appendix G Abbreviations

AoA	Analysis of Alternatives
ARP	Address Resolution Protection
AS&W	Attack sensing and warning
ASR	Address space randomization
AT	Anti-tamper
BCP	Business Continuity Plan
BIA	Business Impact Analysis
BSM	Business Service Management
BYOD	Bring your own device
CAPEC	Common Attack Pattern Enumeration and Classification
CCoA	Cyber course of action
CDS	Cross-domain solution
CERT	Computer Emergency Response Team
CJA	Crown Jewels Analysis
CMP	Chip multiprocessor
CMRS	Continuous monitoring and risk scoring
CND	Computer network defense
CONOPS	Concept of operations
COOP	Continuity of operations plan (or planning)
COTS	Commercial off-the-shelf
CS&IA	Cyber security and information assurance
CSOC	Cyber Security Operations Center
CyOC	Cyber Operations Center
DBMS	Database management system
DCC	Dynamic composable computing
DHCP	Dynamic host configuration protocol
DMZ	Demilitarized zone
DNS	Domain Name Service
DoD	Department of Defense
DOE	Department of Energy
DSR	Data space randomization

EA	Enterprise architecture
FPGA	Field-programmable gate array
GOTS	Government off-the-shelf
HVAC	Heating, ventilation, and air conditioning
I&W	Indications & Warnings
ICT	Information and communications technology
ID	Identifier
IDS	Intrusion detection system
IP	Internet Protocol
ISCP	Information System Contingency Plan
ISR	Instruction set randomization
IT	Information technology
ITSM	IT Service Management
JOC	Joint Operations Center
LOE	Level of effort
MAC	Media Access Control (address)
MIA	Mission Impact Analysis
MILCOM	Military Communications Conference
MLS	Multi-level security
MPSoC	Multiprocessor system-on-a-chip
MSL	Multiple security levels
NIST	National Institute of Standards and Technology
NOSC	Network Operations and Security Center
OS	Operating system
PDA	Personal digital assistant
PM	Program Manager
POET	Political, operational, economic, and technical
QoS	Quality of service
R&D	Research and development
RAdAC	Risk-Adaptable Access Control
RSS	Real Simple Syndication
RTOS	Real-time operating system
SCRM	Supply chain risk management

SED	Self-encrypting drive
SLA	Service level agreement
SME	Subject matter expert
SOA	Service-oriented architecture
SOP	Standard Operating Procedure
SoS	System-of-systems (or systems-of-systems)
TPM	Trusted Platform Module
TRL	Technology Readiness Level
TTP	Tactic, technique, procedure
VMM	Virtual Machine Monitor
WORM	Write-once, read many