Challenges in Applying Resiliency Techniques

Challenges in Applying *Adaptive Response*[1]

Adaptive Response is focused on changing resource configurations and allocations while operations continue seamlessly. In order to make these changes without causing undesired or unanticipated consequences, a potential cyber course of action must be thoroughly tested to ensure consistency with intended results (or at least expected results) before being adopted for use as a potential response. Some of the specific challenges to applying *adaptive response* are:

- 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 resource allocation raises issues of adherence to service level.
- Adaptive Management raises political challenges with respect to responsibilities for ongoing, dynamic risk management. Sacrificing some activities to maintain other activities is best done when risks and priorities are well understood and the authority to make those decisions is clearly allocated. Acquiring this understanding and authority requires a significant amount of planning and practicing.

Challenges in Applying *Analytic Monitoring*

The use of *analytic monitoring* requires cooperation across the constituent systems and organizations in order to identify access to and interactions and dependencies among constituent systems that could indicate access changes, destabilization or disruption before it affects mission performance.

Deliberate efforts are needed to establish monitoring and analysis at the system-of-systems level, to share and fuse information, and to define roles and responsibilities for malware and forensic analysis. Some of the specific challenges to applying *analytic monitoring* are:

- Large volumes of monitoring data create the potential for abuse, especially when large quantities of sensitive data (e.g., privacy data, mission operational planning data) are aggregated together, as the aggregate data set can often be more sensitive than the individual comprising elements. This is especially true of credential related data as the compromise of such data can lead to the compromise of large numbers of accounts.
- A balance between the trade-offs of insights gained from monitoring and the additional protection from encryption, as well as the costs associated with the various options, must be made.
- Coordinating monitoring across architectural layers and across systems as well as across different organizations (with potentially different policies) and for both cyber and non-cyber

[1] All Italicized words are defined in the *Cyber Resiliency Terms and Concepts* Document.

data, is difficult to do. This includes dealing with the lack of visibility into non-owned infrastructures (e.g., networks, cloud computing environments) and data interoperability in fusing and analyzing all the monitoring data. An added complication specific to access control is the multiple identities an individual may possess even within a single organization let alone across several organizations.
- Analytic capabilities, whether in the form of malware analysis, red teaming, or damage assessment, need to keep pace with changes in adversary capabilities as well as in enterprise information and communications technologies. Ongoing investment is needed to meet this need. This is particularly true in the area of baselining typical traffic to isolate and indicate anomalies in traffic that would indicate adversary activity.

Challenges in Applying *Coordinated Defense*

Coordinated Defense requires the coordination of security management, network management, and system management activities in ways that are often not part of the staff job descriptions. These additional responsibilities can create difficulties in obtaining and retaining staff with the needed expertise. Coordinated Defense also requires information sharing which can reveal weaknesses or gaps in an organization's or business unit's governance. Operationally, capturing and presenting information to staff at the level appropriate to their responsibilities, so that they can coordinate and look for inconsistencies, is also a challenges. In addition, each Coordinated Defense approach presents its own challenges:

- Technical Defense-in-Depth uses protective mechanisms, applied at different architectural layers or locations (e.g., using the strategies like those found at https://ics-cert.us-cert.gov/sites/default/files/recommended_practices/Defense_in_Depth_Oct09.pdf). This can increase the cost of development and testing and the complexity of management, training and maintenance.
- Coordination and Consistency Analysis applies processes, supported by analytic tools (e.g., Unified IdAM Administration tools), to ensure that defenses are applied and cyber courses of action are defined and executed in a coordinated, consistent and non-disruptive way. This requires that policy conflicts across organizations be resolved while still respecting equities, particularly as mission and business needs change over time. Changes to component systems and network segments need to be analyzed and tested through exercises that include disruptions in order to identify functional or mission dependencies that may not be known. These conflicts must be resolved and changes implemented in a timely enough fashion that the response is relevant to the ongoing attack.

Challenges in Applying *Deception*

For *Deception* to be effective it is imperative that the adversary believe that the deceptive information or environment is real. Deception can make the adversary uncertain how to proceed, delay the effect of the adversary's attack and increase the likelihood that the adversary will expose tradecraft and TTPs but only if the adversary is fooled by the deception. This requires effective techniques in hiding the real data and providing realistic fake or misleading data and environments. Each of the *Deception* techniques has its own challenge:

- Obfuscation – hiding, transforming or otherwise obscuring information from the adversary – relies on tools such as data encryption. For this to work the encryption keys and other data used to hide the data must be effectively and reliably protected from access by the adversary.
- Dissimulation/Disinformation – providing deliberately misleading information to queries that are suspected of coming from an adversary – poses three main challenges. The first is determining what dissimulation is appropriate, particularly when confused-but-legitimate users could make suspect queries. The second is the operational and economic challenge posed by the significant ongoing effort needed to maintain this approach over time. The third challenge is the audit and legal ramifications of dissimulationand misinformation which can provide challenges to auidts and be potential legal and regulatory liabilities.
- Misdirection/Simulation – maintaining deception resources or environments and directing adversary activities to these resources – presents political, operational and economic challenges. The organization must be willing to commit to supporting this deception environment and resources at a level that is needed for effectiveness. Operationally, the deception environment and resources must be maintained so they appear to be realistic. This is particularly challenging when elements of *dynamic positioning* and *unpredictability* are incorporated into deception environments. There is also the operational challenge inherent in the situations where the adversary executes a multi-pronged attack and not all the prongs are confined to the deception environment.

Challenges in Applying *Diversity*

The use of *diversity* runs counter to the organizational policies requiring adherence to an enterprise architecture including restrictions to a specific set of software products. Operationally, maintaining an accurate representation and consistent management of enterprise systems becomes more challenging as diversity increases. Maintaining IT and help desk support also become more challenging. In addition, for niche products or appliances with narrow functionality, there may be no equivalencies with which to implement diversity. While standards and specifications help, validating that the diverse implementations adhere to these standards and specifications is vital to ensure interoperability to ensure consistency across security and resiliency mechanisms. Some of the specific challenges to applying diversity are:

- Interface standards used in architectural diversity and path diversity, particularly those used for transforming message and data formats must be defined and used consistently. Architectural diversity is often an unintended consequence of systems being acquired at different times by different organizations or to meet different mission needs. This unintended diversity can be a weakness instead of a strength if not managed carefully.
- The different designs used in design diversity must be developed and tested carefully to ensure that they provide equivalent functionality. Design diversity increases the complexity of management, training and maintenance.
- Synthetic diversity can increase the costs of development and testing, to ensure that transformations do not produce inconsistent functionality. For this reason it is important to determine which software elements should be subject to this type of diversity so that the costs are minimized while the impact on the adversary is maximized.

- In order to implement information diversity, mechanisms must be defined and implemented to identify and track the provenance of the information and decide how to handle alternate information based on the level of trustworthiness, validity or quality expected from that source.
- Supply chain diversity can be difficult to verify, as vendors may be unwilling to share – or even unable to identify – sources of components integrated into their products.

Challenges in Applying *Dynamic Positioning*

Dynamic Positioning – distributing and dynamically relocating functionality or assets – requires advanced planning to be effective and not create chaos for defenders. The advanced planning should include how to meet service level agreements when dynamic repositioning is used. This planning should also take into account the need to maintain consistency and integrity for distributed processing and distributed data. In addition there are specific challenges to applying each approach of *Dynamic Positioning*:

- The Functional Relocation of Sensors approach is currently immature. While the mechanisms to implement it are mature, applying it to cyber resiliency is challenging technically and operationally with respect to agility, ensuring synergies with other techniques and handling recovery and evolution.
- The political, operational and technical challenges for the Functional Relocation of Cyber Assets approach largely relate to the transitional status of moving target defense tools. As these tools mature and stabilize, these challenges will decrease.
- The Asset Mobility approach comes from the fields of safety and dependability. The challenges when applying it to cyber resiliency relate to understanding the relationship between physical and logical accessibility.
- The Distributed Functionality approach is extremely mature in many enterprise architectures. The technical challenges in adapting it to cyber resiliency relate to determining what forms of distribution are effective in contested environments – it is critical to consider potential performance impacts as well as limitations due to policy or programmatic restrictions (e.g., organizational commitments to a specific product or product suite which does not accommodate repositioning).

Challenges in Applying *Dynamic Representation*

Dynamic Representation – constructing and maintaining current representations of mission and business posture in light of cyber events and cyber courses of actions – requires trust, or at least a knowledge of the level of trustworthiness, of the data used in constructing the representations. This can be a challenge particularly in relation to non-owned infrastructures (e.g., public cloud computing environments). There are several other specific challenges as well:

- Some of the information used to create a dynamic representation may be sensitive (e.g., mission dependencies, and current adversary characteristics and behaviors) causing the representation to be sensitive. This can also cause the representation itself to become a target of the adversary.
- Not all assessment tools provide information in the same format or at the same level of detail. This can present difficulties when trying to integrate these sources into a single representation.

It is important to identify and document assumptions and decisions made in the integration process.

- The quality of the gathered sensor data can vary, based on sensor capabilities, sensor control, and environment. In addition, the trustworthiness of the data may also be impacted by the adversary trying to deceive the defenders by modifying sensor data in order to create a false picture of the environment.

Challenges in Applying *Non-Persistence*

Non-Persistence – generating and retaining resources as needed or for a limited amount of time – relies on the idea that the particular resource is not needed beyond a certain time. The very nature of non-persistence can be a challenge for certain required processes, particularly those related to cyber defense. For example, non-persistent services could run counter to the need to perform digital forensics to identify the nature of adversary malware, while non-persistent data could run counter to the need to preserve evidence that might be needed for e-discovery, litigation holds or for prosecution. In addition, there are challenges specific to the environment:

- For some business functions, the refresh capability needs to be relatively seamless to ensure that it does not disrupt, or minimizes the disruption of, organizational operations.
- Deletion and sanitization technology of storage media is generally not rapid enough or applied across broad enough spectrums of media to provide high assurance.
- Virtualization is an enabler for achieving non-persistence. However, as not all products and devices support virtualized environments there are limitations to its applicability.

Challenges in Applying *Privilege Restriction*

Privilege restriction requires identifying and resolving the differences between mission and system owners that can result in differences in risk tolerances and trust criteria can differ across component systems. These include inconsistencies or gaps in definitions of roles, responsibilities, and related privileges as well operational impetus to share roles. The use of multiple identifiers across applications, platforms and enterprises can complicate privilege management. In many circumstances, federated identity and privilege management systems can be used to provide needed functionality; however, these may not be useful in highly mobile environments environments (or other environments in which bandwidth or connectivity to such systems is limited). In addition, each *privilege restriction* approach presents its own challenges:

- Privilege Management: mechanisms for identity resolution and/or resolution of other access- or privilege-related attributes across multiple systems are needed; the use of multiple identifiers can complicate privilege management. In many circumstances, federated identity and privilege management systems (e.g., Domain Servers and LDAP servers) can be used to provide needed functionality; however, these may not be useful in environments in which bandwidth or connectivity to such systems is limited.
- Privilege-Based Usage Restrictions: Criteria for usage restrictions that can be applied across component systems need to be defined. There is a danger of lack of agility and flexibility; the mission criticality of a resource can change dynamically and privileges sometimes must change as well in order to ensure the mission.

- Dynamic Privileges: As the mission needs change, least privilege must be maintained at the same time privileges may need to change in order to ensure the mission. Risk-adaptable access control mechanisms (RAdAC) must be identified in order to balance the needs of the mission with the risk posture of least privilege.

Challenges in Applying *Realignment*

Realignment – aligning cyber resources with core aspects of mission/business functions, and thereby minimizing the attack surface – relies on knowing the organization's mission or business functions, knowing (and accepting) their relative priorities, and understanding what aspects are central as opposed to supporting or nice-to-have. Such knowledge – and acceptance of relative priorities – can be politically sensitive within an organization. In addition, each *realignment* approach presents its own challenges:

- Purposing: The trend in enterprise architectures is toward multi-purpose or converged sets of resources, which runs counter to restricting a resource's use to a known set of well-defined purposes. An agile organization that frequently adapts its business model and functions to changing circumstances and new opportunities must make tradeoffs between the cyber resiliency benefits purposing offers and the operational agility it can impede.
- Outsourcing/offloading: Offloading or outsourcing functions to a service provider is more often driven by economics than security. In addition, the organization must ensure that the service provider is capable of providing the needed protection for the inessential functions, and that adequate *segmentation* between the provider's and the organization's systems is implemented.
- Restriction: Removing or disabling unneeded risky functionality or connectivity presents multiple challenges, particularly if the organization depends on commercial-off-the-shelf (COTS) or free-and-open-source (FOSS) software. For COTS software, the organization may be constrained by the terms of its licenses; for FOSS, the organization may lack the requisite in-house technical expertise.
- Replacement: Replacing risky implementations with less-risky ones can be costly, and can be precluded by an organization's commitment to a specific enterprise architecture.

Challenges in Applying *Redundancy*

Redundancy is a highly mature and widely used technique in the area of Contingency Planning, Continuity of Operations, and Performance Optimization. While this is a strength in the stability and wide availability of tools and automation available, it becomes a challenge to modify already existing systems and processes to provide resiliency, as well as address the goals for which the tools were originally intended. In addition, each combination of the *redundancy* approaches presents its own challenges:

- The Protected Backup and Restore approach requires that the information and software must be backed up in a way that protects its confidentiality, integrity, and authenticity, as well as a way to restore it, in case of disruption or destruction, without unnecessary exposure. Appropriately balancing the risk of exposure with the need to provide critical information and software is crucial.
- Surplus Capacity requires accurate analysis to identify functional dependencies and effectively leverage surplus systems.

- Replication likewise, requires accurate analysis for the same reasons.

Challenges in Applying *Segmentation*

The use of *segmentation* runs counter to current trends for integrated services (including integrated communications, enterprise asset management, and Web-enabled shared use of "big data" repositories), convergence of physical and cyber resources, and cloud computing. Unless careful systems engineering is applied, administrative and cyber defender visibility into protected segments within the enterprise may be restricted. In addition, each combination of the segmentation approaches presents its own challenges:

- Predefined physical separation uses redundant hardware and communications media to create physically isolated enclaves. This can increase acquisition, operations, and maintenance costs. In addition, only physical communications media (wired networks) can be physically separated. Operational procedures must be defined for:
    - Validation of storage media (e.g., CDs, USB drives) before such media are used to transfer data or software to the enclave (to avoid, for example, the Stuxnet scenario).
    - Backup and restoration of software, services, and data on the enclave.
    - Disabling wireless communication capabilities provided by devices in physically isolated enclaves. Even with operational procedures in place, users sometimes forget to do this.
- Dynamic physical separation or isolation – changing systems or enclaves while in operation – can involve unplugging devices from networks – for example, removing the cable from the router, switch, or firewall that serves as the gateway between a sub-network and the larger enterprise network. This can have unintended consequences, if the functional dependencies among enterprise services are not well understood. Operational procedures must be defined for making – and, if unintended consequences are intolerable, backing out – changes, in coordination with business process owners.
- Predefined logical separation uses mature technologies – such as encryption, firewalls, virtual machine separation, and access control mechanisms – to restrict the flow of information to, from, or over a device, service, or network, or to and from a data repository. Operational procedures must be defined for correct and effective use of those technologies – for example, changing encryption keys frequently and/or unpredictably, configuring and patching firewalls or other separation technologies, and applying the principle of least privilege to access control.
- Dynamic logical separation uses the same technologies as predefined logical separation, but applies them to reconfigure enterprise resources while trying to avoid interrupting or even degrading service. This requires integration of security information and event management (SIEM) with administration or management capabilities. As with dynamic physical segmentation, operational procedures must be defined for understanding and dealing with the consequences of dynamic changes.

Challenges in Applying *Substantiated Integrity*

For *substantiated integrity* to be effective, it is critical that mission operators and cyber defenders are notified, when threshold conditions are reached, rather than having automated responses to unexpected behavior go unrecognized until a failure occurs. Deliberate efforts are needed to ensure that meta-data is defined and handled consistently. In addition, each combination of the *substantiated integrity* approaches presents its own challenges:

- Integrity/Quality Checks: Integrity checks of data can be done using checksums. Ideally these should be cryptographic checksums to guard against malware modifying the existing checksums or inserting checksums to hide the corruption of the data. Since validating such checksums sometimes is time consuming, it is important that it be done judiciously on critical data and applications. More challenging is the issue of checking the validity of systems and services. For small scale applications the same checksum techniques can be applied. For larger scale, what may be required is polling of inputs from diverse critical services (e.g., Byzantine quorum systems) to determine correct results in case conflicts arise between the services.
- Provenance Tracking requires both trust in the supply chain and the ability to establish the source of the data, software or hardware elements. Both of these can be challenging in today's environment of outsourcing. In addition, this approach may not appropriately identify invalid elements if the outsourced supplier has been successfully attacked by an adversary.
- Behavior Validation requires that baseline behavioral expectations are established. Care must be given when establishing these expectations so that the false positive and false negative rates are appropriately balanced.

Challenges in Applying *Unpredictability*

Unpredictability is not a stand-alone technique; it is used in conjunction with *Adaptive Response*, *Analytic Monitoring*, *Deception*, *Diversity*, *Dynamic Positioning*, *Non-Persistence*, *Privilege Restriction*, and *Segmentation / Isolation*. It must be implemented carefully, to avoid unintended consequences. In particular, unpredictability can present challenges for *Coordinated Defense*. Defenders need to have some way to discern whether an unexpected event is the result of an implementation of unpredictability, or is a possible indicator of adversary activity. In addition there is a significant amount of overhead associated with the creation and maintenance of unpredictability over extended time periods.