A Multidimensional Cyber Threat Scenario Enumeration Model for Resilience Engineering

Anurag Dwivedi

ABSTRACT

Many frameworks have been proposed for analyzing and enhancing the cyber resilience of systems and missions. Most focus on conducting risk or gap analyses before suggesting mitigations. To apply those frameworks, it is essential to gain knowledge about the threat scenarios against which the risk or resilience is being evaluated. Common approaches to threat enumeration include leveraging threat intelligence or identifying sequential actions from threat models that are mainly developed from databases of past threat events. Such approaches either lack comprehensiveness or are too granular to produce a manageable scale of threat action combinatorics when identifying potential cyber threat scenarios for engineering a resilient mission or system. This article suggests a threat scenario characterization and enumeration approach that does not rely on intelligence or past threat databases and allows for tailored abstraction of threat scenarios to inform mitigation strategy decisions and facilitate cybersecurity and resilience engineering.

INTRODUCTION

Although a standardized definition of cyber resilience is still under development, we realized about a decade ago[1,2] that any expression of "resilience" must include (1) a subject with a defined mission or goal (with an identified minimum acceptable performance level to be maintained during a specific period in both normal and stressed operational scenarios) for which the resilience is being described; and (2) a threat or external force against which the resilience is being described. As shown in Figure 1, the subject could be a mission, a system, a subsystem, a device, or a component whose purpose and performance thresholds are well defined. The threat against which the subject's resilience is being explored could be kinetic, natural, nuclear, cyber, or climate related. Resilience can be designed for target threat(s) of a specified type and intensity.

The subject's performance may degrade because of the effects of a threat event, and depending on the resilience mitigations in place, the subject's performance may be fully restored after a period of time. This performance degradation may constitute a failure of the subject unless either the degree of degradation is greater than a unique minimum acceptable performance level or the subject remains in the degraded state for a duration that is shorter than a temporal threshold. Thus, to assess the subject's resilience, one has to know the maximum tolerable bounds for achieving various grades or levels of successes for the subject. For example, a mission may fully succeed, partially succeed if degradation or duration remains within certain bounds, or fail if the performance remains at a certain unacceptable degraded state beyond a threshold tolerance period.

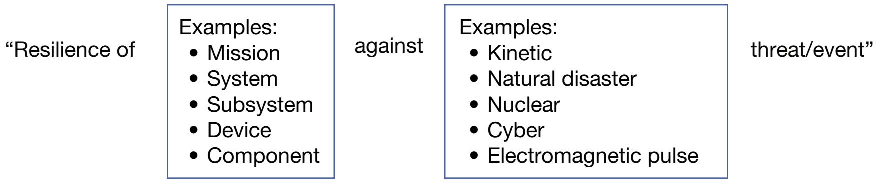

Figure 1. Subject and threat enumeration to properly describe resilience. Any expression of "resilience" must include a subject with a defined mission or goal (and its minimum acceptable performance level over a specified duration) for which the resilience is being described and a threat or external force against which the resilience is being described. This figure illustrates example subjects and threats and their relationship.

Engineering the resilience of a subject thus requires (1) the knowledge of detailed dependencies within the elements of the threat surface of subject; (2) enumeration of relevant threat scenarios at an actionable abstraction; (3) a model for performance degradation (and duration) resulting from compromises caused by the threat; and (4) acceptable performance and duration thresholds for various grades or levels of successes. Obtaining each of these resilience engineering components comes with respective challenges.

This article discusses the challenges associated with identifying relevant threat scenarios and proposes a threat scenario enumeration model (T-SEM) for use in resilience engineering. The T-SEM is abstract enough to cover a comprehensive threat scope and granular enough to inform resilience analysis and mitigation strategy development. Because the threat scenarios capture a full spectrum of threats relevant to resilience engineering and designs, there is no need to enumerate each possible attack vector in extreme detail in early phases of resilience engineering. This limits the enumeration of threat scenarios to a practical scale, ensures that the model is comprehensive and complete in its breadth, and allows for the selection of needed mitigation strategies and approaches starting from early phases of a system's resilience engineering and design life cycle.

RELATED WORK

The term *threat*, in this article, is not defined based on the geography or the threat tier level. Neither is it characterized based on adversary intent, which is generally political, sociocultural, or financial. Rather, the definition of threat scenarios is abstracted to include the attack, target, and access types; the phase of the development cycle; and the broad types of defender's security capabilities targeted by the adversary.

Cyber threat models identify the specific actions that an adversary can take to succeed at each stage toward achieving an offensive malicious goal. The enumeration of threat actions in these models is based on past observations of threat events. However, threat enumeration based on past reported or known incidents does not make up the whole population of relevant threats.

Further, intelligence-based threat forecasts only offer a limited view into future attacks. Factors that can limit these forecasts include intelligence quality, forecast time frame, the intelligence analyst's skill level, reduced visibility into past incidents and future threats, and partial information about adversaries' current or developmental capabilities and intents. As a result, intelligence-centric threat analyses cannot provide a comprehensive cyber threat picture for missions and systems that must be designed to persist for a long lifetime. Resilience designs will likely be insufficient against future threats if they are designed to mitigate only historical threats. Another major challenge in this regard is that the specific start-to-end threat action combinations are too numerous to be practical for individual consideration in resilience engineering and design.

Several threat databases are available, such as MITRE's Common Vulnerabilities and Exposures (CVE),[3] the National Institute of Standards and Technology (NIST) National Vulnerability Database (NVD),[4] the Common Configuration Enumeration (CCE) List (developed by MITRE and transitioned to NIST),[5] MITRE's Common Weakness Enumeration (CWE),[6] and MITRE's Common Attack Pattern Enumeration and Classification (CAPEC).[7]

Threat-based mitigation models include the National Security Agency (NSA)/Central Security Service (CSS) Technical Cyber Threat Framework (NTCTF),[8] the Department of Defense Cybersecurity Analysis and Review (DoDCAR),[9] and MITRE's ATT&CK.[10] Commercial tools can scan a network, network appliance, or element for weaknesses, susceptibilities, or vulnerabilities.[11,12] All of these models depend on databases of vulnerabilities already found or observed in the system and known attack vectors from past incidents.

DEFINING CYBER THREATS AND EXAMPLES OF MITIGATIONS

Security designs based solely on historical incidents are inherently reactive. Since the T-SEM approach discussed in this article does not rely on past events or known vulnerabilities, it captures a comprehensive set of broad threat scenarios for use in security design and mitigation strategy decisions. The T-SEM is based on predefined dimensions of the cyber threats. These include the system's life-cycle phase when threat scenarios and

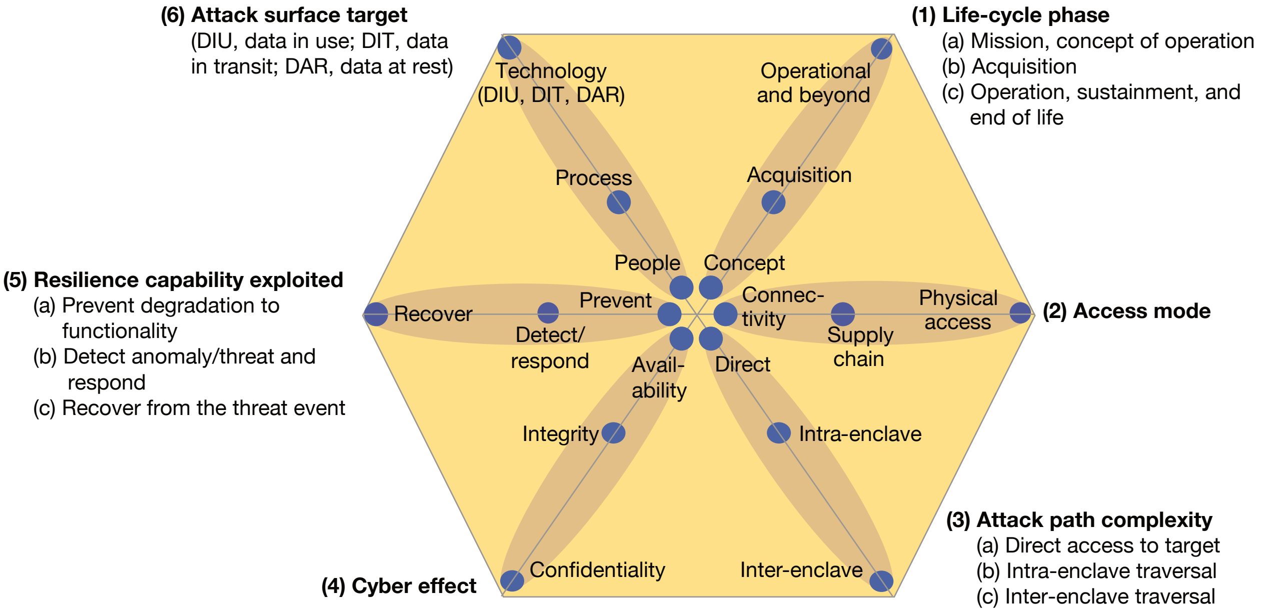
(6) Attack surface target
(DIU, data in use; DIT, data in transit; DAR, data at rest)
(1) Life-cycle phase
(a) Mission, concept of operation
(b) Acquisition
(c) Operation, sustainment, and end of life
Technology (DIU, DIT, DAR)
Operational and beyond
Process
Acquisition
People
Concept
(5) Resilience capability exploited
(a) Prevent degradation to functionality
(b) Detect anomaly/threat and respond
(c) Recover from the threat event
Recover
Prevent
Detect/ respond
Connec- tivity
Physical access
Supply chain
(2) Access mode
Avail- ability
Direct
Integrity
Intra-enclave
(3) Attack path complexity
(a) Direct access to target
(b) Intra-enclave traversal
(c) Inter-enclave traversal
(4) Cyber effect
Confidentiality
Inter-enclave

Figure 2. Multidimensional cyber T-SEM. All dimensions must be considered at each stage of the mission and the system development life cycle to enumerate a comprehensive set of threat scenarios. This spider chart provides only a taxonomy and structure to the threat scenario dimensions and its elements and is not intended to provide a scale along each of its six dimensions.

mitigations are considered, access mode, attack path complexity, compromise or cyber effect type, resilience capability to be exploited, and attack surface node type and data exposure modes. All of these dimensions must be considered at each stage of the mission and the system development life cycle to enumerate a comprehensive set of threat scenarios. These dimensions are displayed in a spider chart in Figure 2, and each is defined later in this section. The spider chart provides only a structure and taxonomy for the threat scenario dimensions and its elements and is not intended to provide a scale along each of its six dimensions.

Five of these dimensions have three elements each, whereas the attack surface target dimension has five elements since the technology element is further divided into those related to data in use (DIU), data in transit (DIT), and data at rest (DAR) sub-elements (discussed in more detail in the section on attack surface nodes). While a maximum of 1,215 combinatorial threat scenarios are possible in this model, not all will be relevant for resilience design for a system supporting a mission in a specific operational environment. The elimination of irrelevant threat scenarios is discussed in the Threat Scenario Enumeration section.

Life-Cycle Phases

Life-cycle phase is a significant dimension to consider from two perspectives: (1) to characterize phase-relevant threat scenarios and (2) to develop phase-specific mitigations for all threat scenarios irrespective of the phase where a threat is invoked. A cyberattack can be planted or launched at any phase of the system development life cycle, including early stages when a mission concept of operations is initially developed. In Figure 3, the system life cycle is characterized by three major phases: pre-acquisition, acquisition, and post-acquisition. To characterize the threat scenario along this dimension, a security engineer would consider only the threats that are relevant and could materialize during any of these three major phases of the system life cycle.[12] Scenarios in the pre-acquisition phase include activities such as

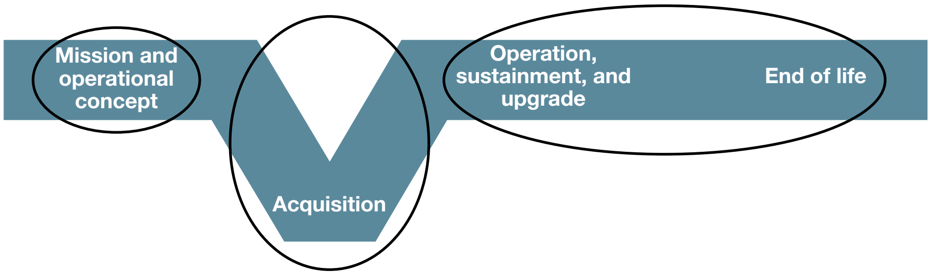
Mission and operational concept
Operation, sustainment, and upgrade
End of life
Acquisition

Figure 3. Life-cycle phases, shown in the systems engineering *V*, where a threat can materialize. Characterization of the threat scenario would consider only relevant threats that could materialize during any of these three major phases. However, mitigations considered in early phases are not limited only to the threats relevant to those early phases but also apply to threats that are planted early but may affect the system or mission in later stages.

development of the mission operational concept. Compromises at this early stage of a program can affect mission operational concepts, weakening resilience, cyber policy, funding for cybersecurity, and the ability to incorporate cyber resilience from the very beginning of the acquisition cycle. Mitigations considered in early phases are not limited only to the threats relevant to those early phases but also apply to those threats that are planted early but may affect the system or mission in later life-cycle stages.

Attacks can be planned, planted, and executed during the acquisition phase and can include compromising not only the systems being acquired but also the acquisition program itself, as well as the infrastructure, processes, and ecosystem used by the program. Affected elements include, but are not limited to, the program protection plan; processes such as requests for information, proposals, or quotations; critical program information; critical technology; supply chain; and other program-originated security and sustainability requirements.

Cyberattacks can happen in the post-acquisition phase during the operational, sustainment, upgrades, maintenance, and end-of-life phases. New susceptibilities can be introduced during maintenance; technology upgrades; improper patches; imperfect operational tactics, techniques, and procedures; or suboptimal implementations of cyber solutions. Because a system's susceptibilities to attack can be exploited during any phase of the life cycle, robust and continuous processes for monitoring, auditing, assessing, and remediating must be required. At the end of its life, a system must be disposed of properly to ensure the confidentiality of sensitive information, technology, vulnerabilities, and processes.

It is worth noting again that the mitigations identified at a particular life-cycle stage are not specific only to the threats relevant to that stage. Mitigations for cyber risks at different stages of the life cycle are unique, justifying the life cycle as an important dimension for characterizing cyber threat scenarios.

Access Mode

An adversary uses three primary modes of access to affect data availability, inject malicious or rogue code, or steal or exfiltrate data or information: (1) external connections to the system via wireless, wired, or other transmission means; (2) rogue code or firmware implanted in the supply chain; and (3) physical access to a cyber system by a human or a robot—for example, when using external media through a USB interface; input/output (I/O) interfaces connecting devices such as a keyboard, video monitor, or mouse (KVM); or a KVM switch. An advanced adversary can gain access by other means, such as modulating a power, acoustic, or laser/optical signal through external I/O interfaces.

The mitigations for each of these three access modes are drastically different. For example, for connected systems, intrusion detection, access control, identity and authorization management, encryption, hashing, allowlist implementation, moving target defenses, and honeypots may be considered. For mitigating supply chain risks, the defender may depend on a trusted supplier or foundry, ensure trusted chain of custody at all times, verify code and validate I/O, implement an allowlist, mandate vendor hash, or use fuzzing-based testing, among other techniques. To mitigate against unauthorized malicious physical access, a defender may control physical and virtual access, protect against burglary, employ hashing, disable unused ports, train users, establish trust, and monitor human behavior and activity trends.

Since the mitigations for safeguarding against these access modes are different, access mode is a justified dimension for characterizing cyber threat scenarios.

Attack Path Type and Complexity

After accessing a system, the adversary can take one of several paths to attack, each with a different level of complexity. The adversary may target a node directly. If direct access is not possible, the adversary may have to traverse through the attack surface topology to access the target. In the latter case, the adversary may have to traverse within a single security enclave or cross security boundaries of multiple enclaves (Figure 4). A directly accessible target may be easier to discover and compromise. Traversal through the attack surface topology nodes may increase the adversary's cost and level of effort. If it is necessary to traverse through multiple security boundaries or enclaves, the adversary's effort, time, and cost may further be elevated.

Mitigations for these three attack scenario elements require astute architecture and design. Critical assets should be behind multiple security boundaries or defenses to limit an adversary's reach and to enhance the defender's ability to detect and respond. The heterogeneity of technology at various nodes along the traversed path will make traversal difficult. Moving target defenses will make the traversal path uncertain for the adversary. A segmentation strategy with multiple enclaves may be employed where enclaves are architected so that critical nodes are not aggregated in a single enclave but rather are distributed over multiple enclaves. A segmentation strategy may enable employment of zero trust security concepts. Additionally, basic protection, detection, preplanned response, and recovery will provide needed cyber hygiene for secure operations.

Compromise Type and Cyber Effect

The type of compromise is another key threat dimension. Cyber compromises manifest themselves in three

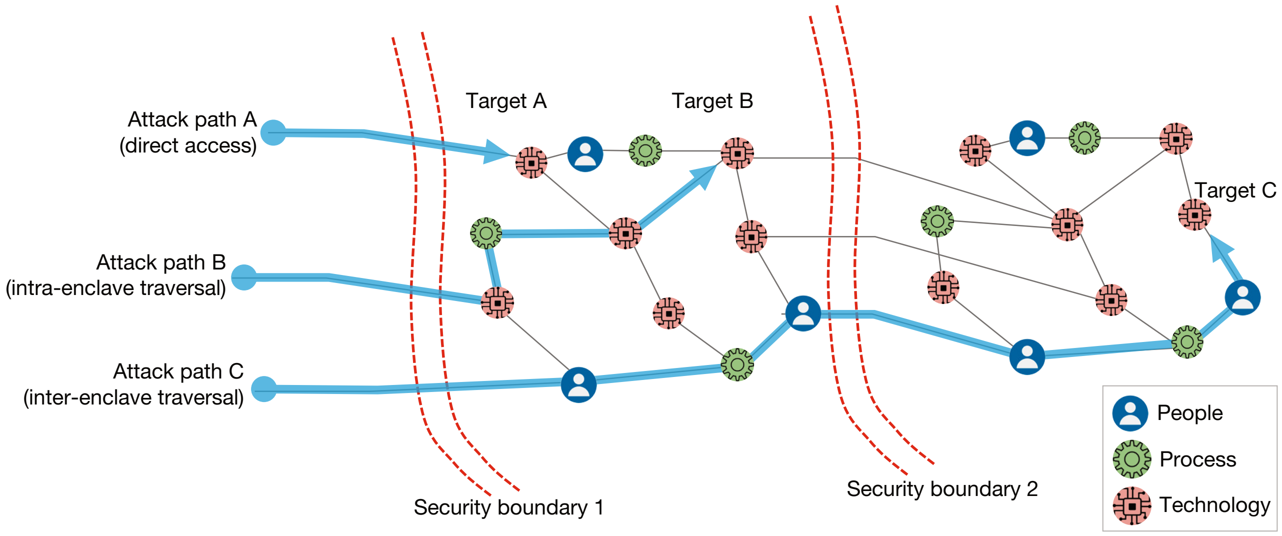

Figure 4. Three types of attack paths to reach the desired targets. Path A is direct access, path B requires intra-enclave traversal, and path C requires inter-enclave traversal. Each path has different costs to the adversary; path C may be the most costly in terms of time and effort.

major ways: confidentiality, integrity, and availability. Confidentiality attacks result in unauthorized exposure or exfiltration of data or information, generally for malicious purposes. Integrity attacks compromise the integrity of data or information. They may result in corrupted data, decreased defender trust in the systems or codes being used, or flawed or degraded system performance. Availability attacks are intended to make data, information, or services unavailable. Examples include disrupting power, creating cyber-physical effects, and denying service.

As illustrated in Figure 5, availability and integrity attacks generally cause more severe tactical loss than confidentiality attacks, whereas confidentiality attacks are responsible for more severe strategic loss. Also, as the figure shows, one type of attack can enable an attacker to succeed subsequently in launching another type. A skillful adversary may use a series of confidentiality attacks during the planning, discovery, or reconnaissance phases of a well-orchestrated cyberattack, with the goal of eventually launching an availability or integrity attack.

It is important to understand the effects of different types of compromises since they will have varying degrees of impact to an organization's, or mission's, tactical and strategic goals. Resilience against these three types of compromises may depend on the nature of the organization's or mission's goals. For example, financial institutions may be able to quantify the impact of a confidentiality attack more readily than the defense sector can, where the loss may be strategic and harder to quantify. Integrity attacks may be more serious for tactical defense missions than confidentiality attacks. Accordingly, mission owners and organizations may prioritize mitigations according to the consequences they may suffer from such compromises.

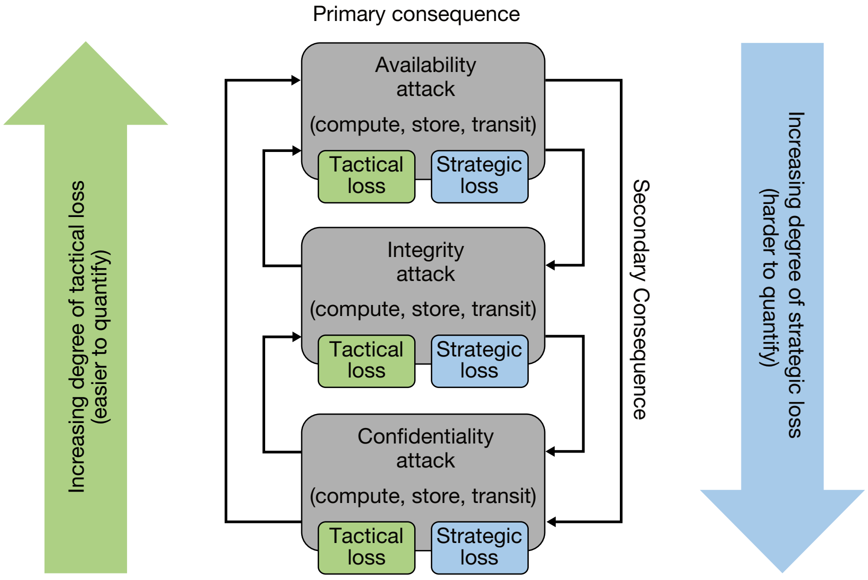

Figure 5. Confidentiality, integrity, and availability compromises with tactical and strategic consequences. While availability and integrity attacks generally cause more severe tactical loss, confidentiality attacks are responsible for more severe strategic loss. Also, one type of attack can enable an attacker to successfully launch another type.

Mitigations for confidentiality compromises are primarily obfuscation, strong encryption, and access control. Integrity can be detected using hash comparison and I/O verification and can be prevented by limiting read/write access and implementing redundancy combined with voting schemes. Availability attacks can be countered with preprovisioned redundancy when supported by heterogeneous technologies and security architectures.[13–17]

Since these compromises require different types of mitigations and have varying degrees of impact to an organization or mission, the compromise type and effect is an important dimension for cyber threat scenario characterization.

Targeted Security and Resilience Aspects

The NIST cybersecurity framework[17] identifies five core functions of cybersecurity: identify, protect, detect, respond, and recover. These map well to the Cyber Survivability Endorsement framework,[18] which describes all these functions using only three survivability aspects: prevent, mitigate, and recover. No matter which framework is used, these functions are essential cybersecurity and resilience design elements, and an adversary may target any of them. In addition to protecting basic functionality, a good cybersecurity and resilience practice is to implement appropriate detection/response and recovery methods. This dimension of threat scenario assumes two security postures: (1) the defender has not properly implemented prevention, detection/response, and recovery defensive controls, and adversary methods simply exploit their absence; and (2) such measures are in place, but adversaries are able to degrade or defeat them to achieve their goals.

Compromising the ability to protect the system may degrade its (or its elements') functionality and may have cascading effects that propagate eventually to affect the mission. An adversary may realize such effects by compromising physical or logical access controls, compromising obfuscation techniques such as encryption, or bypassing allowlists.

An adversary may choose to attack the detectability of malicious activities and anomalous behavior in a cyber system. If the system is designed to sense, detect, alarm, log, and alert to any intrusions, anomalies, or unexpected performance, compromises to these capabilities may have serious consequences or allow an adversary to hide moves or progress the attack along the intended attack vector. Common mitigations against compromises to detection capability are to implement a separate out-of-band, actively monitored detection system with a separate security enclave and to institute privilege access or escalation processes.

Response and recovery capabilities are tightly coupled with detection capabilities. Anomaly detection could trigger an automated or operator-assisted response. This response may be tactical remediation within mission-relevant time frames, even if it degrades system performance. If the tactical remedial response is not sufficient to achieve tactical goals, restoration and recovery may be necessary. Restoration and recovery may or may not be completed in the mission-relevant or desired time frame. Response and restoration capabilities may be protected by frequent checks and audits to ensure that all enabling elements, particularly those providing a backup or redundant capability to a primary means, are in working order and will function as expected when needed.

Mitigations for protection or prevention, detection/response, and recovery capabilities are distinct and may affect the resilience of the mission. For these reasons, adversarial compromise of these resilience measures is a unique and essential dimension of cyber threat scenario characterization.

Attack Surface Node Type and Data Exposure Mode

Cyberattack surface can include people, processes, and technology elements that can be identified from a comprehensively described system model. While the system model may contain many systems, subsystems, components, interfaces, data and service flows, operators, processes, and procedures, the cyberattack surface may comprise only a small subset of those entities. Cyberattack surface constitutes only a subset of the complete system model and includes only those elements that are cyber relevant. Cyberattack surface enumeration, however, must consider both the internal and external cyber-relevant entities if they have common service interfaces. In this article, the elements of the cyberattack surface are called the nodes of the attack surface graph. The number of nodes scales consistently with the abstraction level of the attack surface description or topology.[19]

A threat can target a people node, a process node, a technology node, or a combination[2,19] to eventually compromise data or services. Depending on the relationships between the attack surface nodes and subject performance, analysis can assess the impact[19] of a compromise. Also, depending on the node's contribution to the subject's performance, with or without response and restoration, its criticality[20] can be determined. Criticality analysis does not require knowledge of a detailed attack vector since it considers the mission impact *if* (and not *how*) a node is compromised to achieve an intended cyber effect. In one use case, criticality analysis can prioritize the application of mitigations among the attack surface nodes based on their relative criticalities. This may limit the complexity and cost of applying appropriate mitigations, while meeting resilience design goals.

Mitigations for people, process, and technology nodes may be different, justifying the need for identifying the attack surface node type as a dimension to characterize the threat scenario. Mitigations for people and process nodes can be envisioned to be governance centric, zero trust, and implementing dynamic access policies based on continuous advanced analytics such as behavioral factors, machine learning, and artificial intelligence.

Technology nodes include three primary modes[2] in which data can be exposed: DIU, DIT, and DAR. Data processed by a computing device are referred to as DIU. In most applications, even if the data are transported and stored encrypted, they are decrypted for computational processing. Thus, DIU is a critical exposure mode available to a cyber adversary. Research on and early applications of homomorphic encryption may alleviate or obviate the need to decrypt while processing. However, several currently available homomorphic encryption solutions require significant computational resources, which may negatively affect system performance. While these technologies continue to mature, mitigations for DIU compromise remain indirect, such as those enabled through access control.

Data stored on disks, on removable media, and in databases are referred to as DAR and provide another exposure mode to cyber adversaries. Mitigations such as heterogeneous redundancy, full-disk encryption, and access control protect DAR against confidentiality, availability, and integrity attacks.

Data transmission within a system or between systems using local or wide area networks or direct (wired or wireless) communications links provides another exposure mode to the cyber adversary. The complex layering and protocols involved in data transport present numerous opportunities for a cyber adversary to compromise DIT. For example, protocols at each Open Systems Interconnection (OSI) layer, encapsulations, and tunneling may provide multiple levels of exposure and compromise opportunities. Encrypted data may have to be decrypted at the transit nodes. In some cases, only the header, and not the payload, needs to be decrypted to facilitate transport and routing. Both the header and payload may be allowed to be encrypted for encapsulated and tunneled data for a specific application. Because of the complexities involved in end-to-end secure data transmission, defense-in-depth must be carefully applied to protect against all exposures for multi-layered, multi-protocol, multi-encapsulated transmission. Improperly secured transmission may allow attacks such as man-in-the-middle and replay, among other malicious activities.

Mitigation strategies for each of these three exposure modes may differ significantly, justifying the DIU, DAR, and DIT exposure modes as unique sub-elements of technology nodes as a threat scenario attack surface target element.

THREAT SCENARIO ENUMERATION

Adversarial threat action paths along the cyber kill sequence can be determined by granular threat models such as ATT&CK. Representing all combinations of threat actions in the ATT&CK model would result in many billions of combinations, making development of a mitigation strategy extremely challenging because of the sheer scale of attack scenarios requiring mitigations. Security engineers use multiple approaches to suppress this prolific set of attack path scenarios, such as random selection, Monte Carlo–assisted selection, selection of a subset of threat path scenarios, or use of intelligence information to identify the most likely threat path scenarios.

While any of these approaches may produce a smaller and more manageable threat path set, they are all insufficient for risk assessment and resilience engineering. If mitigations are implemented for a randomly selected small subset of threats, adversaries will identify a non-mitigated path in the early reconnaissance phase of the cyber kill sequence. Monte Carlo selection will not

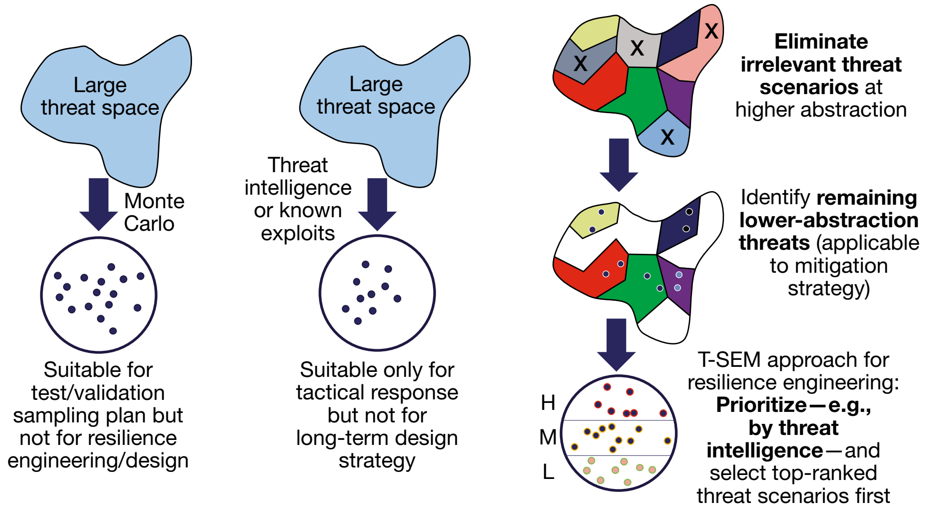
Large threat space
Monte Carlo
Suitable for test/validation sampling plan but not for resilience engineering/design
Large threat space
Threat intelligence or known exploits
Suitable only for tactical response but not for long-term design strategy
X
X
X
X
X
Eliminate irrelevant threat scenarios at higher abstraction
Identify remaining lower-abstraction threats (applicable to mitigation strategy)
H
M
L
T-SEM approach for resilience engineering: Prioritize—e.g., by threat intelligence—and select top-ranked threat scenarios first

Figure 6. Suppression of enumerated threat scenarios. Rather than relying on Monte Carlo selection or intelligence-based selection, the T-SEM abstracts threat scenario enumeration by looking at six distinct dimensions. Elements that are not relevant can then be eliminated, and remaining elements can be further decomposed and then prioritized based on, for example, threat intelligence or threat criticality.

result in selection of the most likely or critical threat paths to be mitigated because cyberattacks are deliberate, and not proven to be random or stochastic. As with random selection, if a small subset of threats is chosen, adversaries will identify a nonmitigated path. Threat intelligence is usually more useful for short-term tactical response and defensive posture. While it may be useful for short-term prioritization of threats to mitigate, short-term threat intelligence should not be used for long-term resilience design.

The T-SEM approach suggested in this article is an abstracted threat scenario enumeration that covers a large threat space with six distinct dimensions. Each dimension has three elements (except for the attack surface component, which effectively has five elements, namely people; process; and DIT, DIU, and DAR technology targets). This yields a total combinatorics of 1,215 threat scenarios, which is many orders of magnitude smaller than the ~10^{18} threat path scenarios computed from a more granular ATT&CK model. The T-SEM threat scenarios of interest, however, can be further reduced by eliminating dimensional elements that are not relevant to the mission, system, or operating environment, as illustrated in Figure 6. Remaining threat dimensional elements or scenarios can be prioritized using criteria including, but not limited to, threat intelligence.

Identification of relevant dimensions and threat taxonomy for characterizing cyber threats allows enumeration of the relevant threat scenarios. The proposed multidimensional model shown in Figure 2 provides a basis for enumerating threat events at an abstraction level sufficient to identify and mitigate gaps in the early concepts of operation, architecture, governance, security policy, and high-level design. Threat scenario enumeration is also helpful in facilitating mitigation trade space and decision analyses. The abstracted threat scenarios ensure completeness of threat coverage while managing the scale and size of threat scenario enumeration.

Figure 7 illustrates an example of suppressing threat scenarios for a specific mission, system, and operational environment. The system is in its operational and sustainment phase of the life cycle, where an adversary can access the system only through its connections with a larger network and supply chain (and cannot gain physical access) because physical interfaces are either removed or robustly protected. This system has only a single security enclave, and critical nodes are behind a gateway requiring an intra-enclave traversal to reach the targeted critical system nodes. Because of the nature of the mission and its dependence on the system, the mission owner is only concerned about availability attacks. The defender wants to consider proper prevention, detection/response, and recovery aspects to achieve the desired resilience but

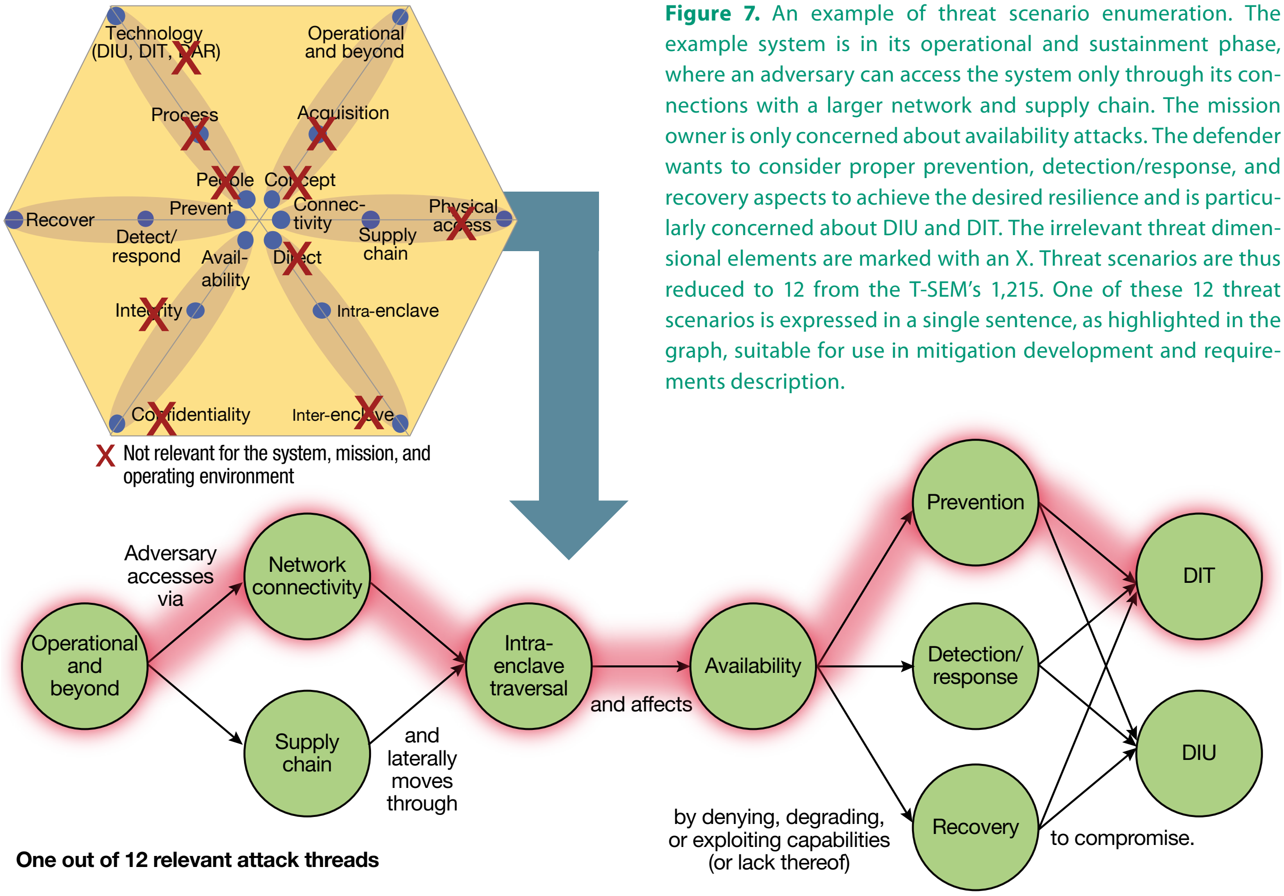

Figure 7. An example of threat scenario enumeration. The example system is in its operational and sustainment phase, where an adversary can access the system only through its connections with a larger network and supply chain. The mission owner is only concerned about availability attacks. The defender wants to consider proper prevention, detection/response, and recovery aspects to achieve the desired resilience and is particularly concerned about DIU and DIT. The irrelevant threat dimensional elements are marked with an X. Threat scenarios are thus reduced to 12 from the T-SEM's 1,215. One of these 12 threat scenarios is expressed in a single sentence, as highlighted in the graph, suitable for use in mitigation development and requirements description.

is not sure whether these resilience aspects are properly implemented. Also, the defender is particularly concerned about DIU and DIT; since there are no critical data stored for accomplishing the mission, DAR is not relevant in this example. People and processes are trusted and are not considered key targets in relevant threat scenarios.

In the figure, the irrelevant threat dimensional elements are marked with an X. Threat scenarios are thus reduced to 12 from the T-SEM's 1,215. One of these 12 threat scenarios is expressed in a single sentence suitable for use in a mitigation requirements description. For such a system, if there are *N* critical nodes to be targeted, there are 12*N* potential threat scenarios. This example demonstrates that while a comprehensive threat scenario-based analysis may be perceived to be difficult, it is well within the realm of practical implementation of deterministic risk analysis and mitigation strategy. Mission resilience against each of these threat scenarios can be analyzed and a mitigation strategy can be developed.

Once this enumeration of relevant threat scenarios is complete, the list can further be reduced or prioritized by applying the information specific to the operating environment or if definitive likelihoods or prioritization criteria are known from threat intelligence. Specific attack vectors for relevant T-SEM threat scenarios can be developed using more granular cyber kill sequence models such as ATT&CK to support the identification of specific cyber solutions aligned with the mitigation strategy.

APPLICATION OF T-SEM TO RESILIENCE ENGINEERING

Key applications of threat scenarios are in assessing, engineering, designing, and enhancing cyber resilience and ensuring mission survivability. While the intent of this article is not to discuss exhaustive use cases of threat characterization and enumeration, an example use case, a simple process for resilience evaluation, engineering, and design, is illustrated in Figure 8.

In this example, a description of cyber-relevant system nodes (e.g., through the cyberattack surface enumeration process[19]) is needed, along with a description of the data or services provided by those nodes as well as activities relevant to other system nodes or supporting essential mission process steps. The dependencies between system nodes and mission functions allow an assessment of mission performance or impact degradation when a specific threat scenario, involving a specific cyber compromise at a specific system node, materializes. Proper mission engineering builds mission resilience through contingencies at the operational level to guard against system function degradations or failures. A key step in the development of a prioritized mitigation strategy is identifying (1) specific threats—a combination of node(s) and compromise(s)—that are capable of degrading mission performance below its tolerance threshold and (2) whether response or restoration can revert the system to an acceptable level of mission performance

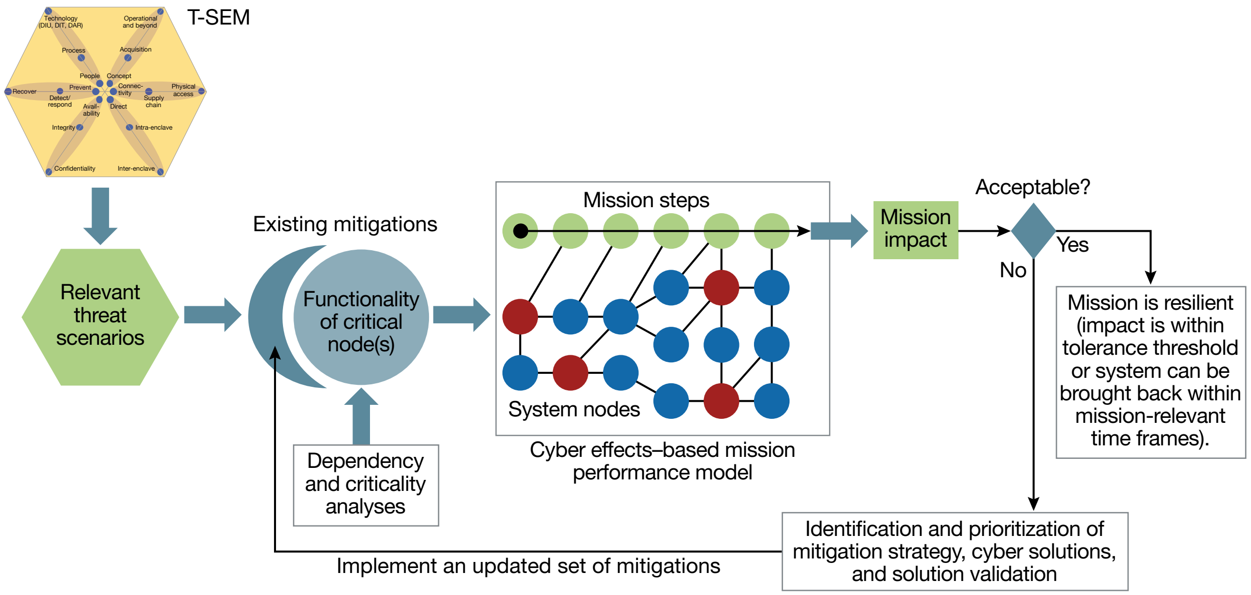
T-SEM
Technology (DIU, DIT, DAR)
Operational and beyond
Process
Acquisition
People
Concept
Recover
Detect/ respond
Prevent
Connec- tivity
Supply chain
Physical access
Avail- ability
Direct
Integrity
Intra-enclave
Confidentiality
Inter-enclave
Relevant threat scenarios
Existing mitigations
Functionality of critical node(s)
Dependency and criticality analyses
Mission steps
System nodes
Cyber effects–based mission performance model
Mission impact
Acceptable?
Yes
No
Mission is resilient (impact is within tolerance threshold or system can be brought back within mission-relevant time frames).
Identification and prioritization of mitigation strategy, cyber solutions, and solution validation
Implement an updated set of mitigations

Figure 8. Role of threat scenario enumeration in resilience design. In this example, a description of cyber-relevant system nodes is needed, along with a description of the data or services provided by those nodes as well as intra-node activities that provide services to other system nodes or to essential mission functions or process steps. The dependencies between system nodes and mission functions allow an assessment of performance when a specific threat scenario materializes. If the resilience is not sufficient, mitigation approaches need to be identified and implemented, and the analysis can be repeated to assess whether the enhancements are sufficient. The process can be iterated until the desired resilience is achieved against the design threat scenarios and threat intensities.

within mission-relevant time frames. Detailed criticality analysis methodology is published elsewhere.[20]

If the resilience of the starting system architecture or design is not sufficient, mitigation approaches need to be identified and implemented.[2,19–21] The analysis shown in Figure 8 can be repeated after considering new mitigations to assess whether the resilience enhancements meet the challenges of relevant threat scenarios. The process can be iterated until the desired resilience is achieved against the design threat scenarios and threat intensities.

As a defender, a consequence-aware mitigation strategy can be used for security and resilience engineering and design processes. This will allow the defender to understand what nodes or combination of nodes need to be hardened against which type of cyber compromises to achieve the desired level of resilience, without requiring the defender to apply all mitigations uniformly to all nodes.[20]

Figure 9a displays the entire threat scenario space from the T-SEM schematically. Not all of these 1,215 scenarios are relevant for every system, mission, and operational environment, so applicable node-specific threat scenarios need to be identified as a reduced set, as shown in Figure 9b, for a specific node of the system.

Once the critical nodes and respective most critical threat scenarios (those that have the potential for the most severe mission impact) are identified, mitigation strategies for each of the threat scenarios associated with each critical node can be astutely determined by cybersecurity subject-matter experts. When developing mitigation strategies, these experts would consider threat coverage, mitigation effectiveness, affordability, feasibility, and practicality, among other factors. Mitigation approaches identified for each of the threat scenarios can be added to a mitigation database as a resilience design utility for future use. If the mitigations applied to each node against each threat dimensional element are added to a database, post-mitigation threat coverage (Figure 9c) can be visualized using a suitable visualization tool.

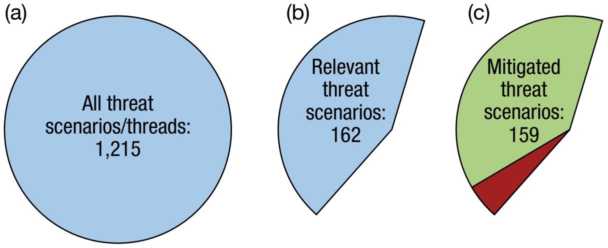
(a)
All threat scenarios/threads: 1,215
(b)
Relevant threat scenarios: 162
(c)
Mitigated threat scenarios: 159

Figure 9. Visualizing relevant threat scenarios, and their coverage by mitigations, to assist resilience engineering. (a) The entire threat scenario space from the T-SEM includes 1,215 possibilities. (b) Applicable node-specific threat scenarios need to be identified as a reduced set for a specific node of the system. (c) Mitigation approaches identified for each of the threat elements or scenarios can be added to a database, and post-mitigation threat coverage can be visualized using a visualization tool.

Mitigation approaches can be categorized as architecture (mission, system, and security), technology (cyber solutions), and governance approaches.[2] A structured approach to identifying and prioritizing mitigations is essential for resilience design and engineering (Figure 9c) for any system but is outside the scope of this article. Mitigation approaches and their implementation must be affordable and feasible, considering the system, application, use case, and constraints of the operating environment. Optimizing any mitigation strategy requires knowledge of critical combinations of threats and system nodes, which is where the abstracted T-SEM-based thread enumeration is useful. This approach comprehensively identifies critical areas by abstracting and then enables identification of specific cyber solutions by drilling down on specific critical areas.

KEY CHALLENGES AND FUTURE WORK

While this article presents a comprehensive and manageable T-SEM, mitigation strategy development for each of the threat scenarios is not within its scope. When suppression of threat scenarios is considered, some of the combinations may not make sense for any system, mission, or operational environment. These are generally related to the combinatorics of people and process nodes with typical cyber threat elements. A follow-on effort could identify those combinations and eliminate them from the maximum possible 1,215 threat scenarios. An optimal set of global or node- or element-specific mitigations will obviate the need for mitigating each and every element of a six-layer T-SEM individually. These advanced, structured, and efficient mitigation identification, prioritization, and validation schemes could be developed for implementing efficient and affordable resilience engineering. Finally, follow-on work could more fully validate the hypothesis that the T-SEM offers a comprehensive description of all applicable threat scenarios and offers value and effectiveness in resilience engineering.

REFERENCES

[1] A. Dwivedi, "Quantifying resilience," tutorial presented at the 5th Ann. Workshop on Cyber Resilience by MITRE, May 20-21, 2015.

[2] A. Dwivedi, "Designing for resilience," *Proc. SPIE 9097*, Cyber Sensing 2014, 90970C-9, Jun. 18, 2014, https://doi.org/10.1117/12.2054389.

[3] CVE (Common Vulnerabilities and Exposure), MITRE, https://cve.mitre.org (accessed Dec. 29, 2022).

[4] NVD (National Vulnerability Database), NIST, https://nvd.nist.gov (accessed Dec. 29, 2022).

[5] CCE (Common Configuration Enumeration). MITRE,. https://cce.mitre.org (accessed Dec. 29, 2022).

[6] CWE (Common Weakness Enumeration). MITRE. https://cwe.mitre.org (accessed Dec. 29, 2022).

[7] CAPEC (Common Attack Pattern Enumeration and Classification). MITRE. https://capec.mitre.org (accessed Dec. 29, 2022).

[8] NSA Cybersecurity Operations, "The Cybersecurity Products and Sharing Division, "NSA/CSS Technical Cyber Threat Framework v2," Nov 29, 2018, https://media.defense.gov/2019/Jul/16/2002158108/-1/-1/0/CTR_NSA-CSS-TECHNICAL-CYBER-THREAT-FRAMEWORK_V2.PDF.

[9] P. Dinsmore, "DoDCAR/.govCAR," presented at the Software and Supply Chain Assurance Forum (SSCA), McLean, VA, Sep. 26–27, 2018, https://csrc.nist.gov/CSRC/media/Projects/cyber-supply-chain-risk-management/documents/SSCA/Fall_2018/WedPM2.2-STARCAR%20SCRM%20FINAL%20508.pdf.

[10] ATT&CK. MITRE. https://attack.mitre.org (accessed Dec. 29, 2022).

[11] OWASP. "Vulnerability scanning tools." https://owasp.org/www-community/Vulnerability_Scanning_Tools (accessed Dec. 29, 2022).

[12] Department of Defense, "DoD Program Manager's Guidebook for Integrating the Cybersecurity Risk Management Framework (RMF) into the System Acquisition Lifecycle," Ver. 1.0, Sep. 2015, https://www.dau.edu/tools/t/DoD-Program-Manager-Guidebook-for-Integrating-the-Cybersecurity-Risk-Management-Framework-(RMF)-into-the-System-Acquisition-Lifecycle.

[13] "NIST Risk Management Framework." NIST. https://csrc.nist.gov/projects/risk-management/risk-management-framework-(RMF)-Overview (accessed Dec. 29, 2022).

[14] NIST, "Security and privacy controls for federal information systems and organizations," NIST Special Publication 800-53, Rev. 5, https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final.

[15] M. Reed, "Vulnerability analysis techniques to support trusted systems and networks (TSN) analysis," presented at the 17th Ann. NDIA Sys. Eng. Conf., Springfield, VA, Oct. 29, 2014, https://ndiastorage.blob.core.usgovcloudapi.net/ndia/2014/system/16997WedTrack1Reed.pdf.

[16] D. J. Bodeau and R. Graubart. "Cyber Resiliency Engineering Framework," MITRE Tech. Rep. MTR110237, Sep. 2011, https://www.mitre.org/news-insights/publication/cyber-resiliency-engineering-framework.

[17] NIST, "Framework for improving critical infrastructure cybersecurity," ver. 1.1, NIST, Apr. 16, 2018, https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf.

[18] D. Clothier, "New DoD approaches on cyber survivability of weapon systems," presented at the AFCEA Chapter Event (Alamo ACE) 2016, San Antonio, TX, Dec. 5–8, 2016, http://www.alamoace.org/resource/resmgr/2016_ace/2016_speakers/doc_clothier_dean.pdf.

[19] A. Dwivedi, "Implementing cyber resilient designs through graph analytics assisted model based systems engineering," in *Proc. IEEE Int. Conf. on Softw. Qual. Rel. and Secur. Companion (QRS-C)*, 2018, pp. 607–616, https://doi.org/10.1109/QRS-C.2018.00106.

[20] A. Dwivedi and D. Tebben, "Cyber situational awareness and differential hardening," in *Proc. SPIE. 8408, Cyber Sensing 2012*, 840803, 2012, https://doi.org/10.1117/12.915642.

[21] A. Dwivedi, D. Tebben, and P. Harshavardhana, "Characterizing cyber-resiliency," *Proc. 2010 Mil. Commun. Conf. (MILCOM)*, 2010, pp. 1304–1309, https://doi.org/10.1109/MILCOM.2010.5680128.

Anurag Dwivedi, Asymmetric Operations Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Anurag Dwivedi is a cyber systems security engineer in APL's Asymmetric Operations Sector. He has a BS in materials engineering from the Institute of Technology, Varanasi, India, an MS and a PhD in materials science and engineering from Alfred University, and an MS in information systems from Johns Hopkins University. Anurag has a 30 years of experience in research, project management, and modeling and simulation, as well as technical leadership, supervision, training, and coaching. He is a recognized expert in cyber and critical infrastructure resilience and security, systems engineering of enterprise and embedded tactical systems, model-based and secure cyber systems engineering, zero trust, and cloud and information security. He has extensive knowledge of standards, Risk Management Framework and conformance, graph-based cyber resilience analytics, cyber network planning and design, communications traffic estimation and forecasting, network applications, airborne networks, optical fiber, fiber and free-space optic systems, and directional mobile ad hoc network (MANET). Anurag has published and presented extensively on these topics. His email address is anurag.dwivedi@jhuapl.edu.