

Protecting Mission Critical Systems The Need for a Shift in Culture, Strategy, and Process

Ron Ross**, ron@ronrossecure.com; Kymie Tan, kymie.tan@jpl.nasa.gov

Copyright ©2025 by Ron Ross and Kymie Tan. Permission granted to INCOSE to publish and use.

** Former Fellow, National Institute of Standards and Technology

■ ABSTRACT

In contrast to the traditional compliance-based approach to protecting space systems using the NIST Risk Management Framework (RMF), a trustworthy secure systems engineering approach as described in the NIST Special Publication 800-160 is proposed as a viable and effective alternative. This paper discusses the issues and concerns with the traditional approach to cybersecurity and how engineering-based approaches measurably improve security, allowing a greater return on investment for mission critical operational environments like those that support space missions. The paper will show that there are several facets to the cybersecurity problem that go beyond the technical to include culture, process, and policy, and explain why a change in strategy and approach is necessary to address the modern sophisticated cyber adversary operating in a world of highly complex and evolving systems. Insights from a project where a NIST SP 800-160-based engineering approach was applied to secure a space mission will be discussed. The early lessons not only illuminate the benefits of security systems engineering, but also the effect of culture, policy and process on building resilience into mission critical systems.

■ **KEYWORDS:** trustworthy secure systems; secure-by-design; systems security engineering; cyber-resilient systems; securing space systems; assurance; systems engineering; security design principles; advanced persistent threat; authorization-to-operate; mission risk; system life cycle

INTRODUCTION

Space is an essential component to the modern economy and vital to the national and economic security interests of the United States. The space sector is critical to many industries, including telecommunications, navigation, and the defense industrial base. Engineering trustworthy, secure space systems is a significant undertaking that requires a substantial investment in the requirements, architecture, and design of systems, components, applications, and networks. A trustworthy secure space system is engineered to provide compelling evidence to support claims that it meets its stakeholder requirements to deliver the capability, protection, and performance needed by the organizations investing in the technology. Adopting a disciplined, structured, and standards-based set of systems security

engineering activities and tasks provides an important starting point and forcing function to initiate a needed change toward defensible space systems that are resilient to the modern adversary.

Building trustworthy, secure space systems cannot occur in a vacuum with “stovepipes” for software, hardware, information technology, and the human element (e.g., designers, operators, users, and adversaries of these systems). Rather, it requires a transdisciplinary approach to protection, a determination across all assets where loss could occur, and an understanding of adversity, including how adversaries attack and compromise systems. This paper addresses considerations for the engineering-driven actions necessary to develop defensible and survivable space systems, including the components that compose,

and the services that depend, on those systems. The objective is to address security and resilience issues from the perspective of stakeholder requirements and protection needs and to use established engineering processes to ensure that such requirements and needs are addressed with appropriate fidelity and rigor across the entire life cycle of the system.

BACKGROUND

In 2002, the United States Congress passed the Federal Information Security Management Act (FISMA) (Anon. 2014), affirming the government’s commitment to protecting the confidentiality, integrity, and availability of federal information and information systems. As part of the FISMA legislation, the National Institute of Standards and Technology (NIST), a bureau

within the Department of Commerce, was given important responsibilities for developing and implementing cybersecurity standards and guidelines for the federal government and its contractors to ensure compliance with the law. In fulfillment of its FISMA responsibilities, NIST developed the Risk Management Framework (RMF) (Joint Task Force (JTF) 2018) and a series of supporting standards and guidelines to help organizations build, operate, and continuously monitor their information security programs. The publications included standards for security categorization (NIST 2004) and minimum-security requirements (NIST 2006), a comprehensive catalog of security and privacy controls (JTF 2020a), and detailed assessment procedures (Joint Task Force Transformation Initiative 2022) to determine if the controls were implemented correctly, operating as intended, and producing the desired effect with regard to enforcing the organization's security policy.

In accordance with FISMA and the Office of Management and Budget (OMB) policy (OMB 2016), the heads of federal agencies were responsible for managing the information security risks associated with operating their information systems. The NIST RMF was the primary vehicle used by agencies to protect the information being processed, stored, and transmitted by their systems. Every federal information system was required to receive an authorization to operate (ATO) prior to being deployed into operational environments to carry out federal agency missions and essential functions. The ATOs had to be signed by the heads of the respective federal agencies or their designated representatives. The ATOs conveyed the information security risk accepted by senior leaders after they had implemented all of the required safeguards and countermeasures (i.e., security controls) needed to protect their information and information systems.

THE PROBLEM

The Risk Management Framework and its supporting publications were designed largely for enterprise information technology (IT) systems. These systems, for the most part, were composed of commercial off-the-shelf hardware, software, and firmware components. This has been the primary focal point for the RMF since its inception in 2005. In subsequent years, the framework and controls were applied to operational technology (OT) systems and IoT devices. While the RMF has been effective in the context for which it was designed, it has been less effective when applied to large and complex systems engineering efforts, for example, in DoD weapons systems and

the NASA's space systems. This problem has been exacerbated by the convergence of cyber and physical systems and the emergence of artificial intelligence (AI) and robotics technologies. In addition to the above, cybersecurity has largely been implemented as a separate and disconnected process for the past four decades creating several institutional and generations problems. These include:

- Insufficient alignment with the systems engineering life cycle of complex systems, creating a disconnected process
- Insufficient attention to risks involving cyber-physical assets (e.g., application specific integrated circuits, FPGAs, programmable logic controllers, robotic actuators, sensors)
- Inadequate integration of cybersecurity risks into the established framework for overall project risks (e.g., safety, reliability)
- Inadequate conversion of current threat intelligence into actionable items by systems engineers
- Questionable protection, ambiguous return on investment (e.g., unknown confidence or assurance against a range of specified threats)
- Inadequate visibility into the underlying system design resulting in insufficient trust and assurance in the system capability
- Ineffective for emerging technologies like AI, autonomy, and cloud-based ground stations, insufficient guidance is provided on how to secure these cutting-edge systems effectively or in a timely fashion.

To address these problems, NIST developed a set of systems security engineering (SSE) tools and approaches to help organizations developing systems for their critical missions. The SSE guidance is contained in NIST SP 800-160, Vols. 1 and 2 (Ross, Winstead, and McEvilly 2022; and Ross, et al. 2021). The engineering-based security approach was designed to help organizations address their protection needs for complex systems, manage the risk of uncertainty during the development process, and provide sufficient evidence to authorizing officials to make informed, risk-based decisions on approving systems for operation. However, despite the comprehensive NIST guidance, organizations have been reluctant to adopt the engineering-based security approach to satisfy FISMA and OMB security compliance requirements. The next sections provide additional details on the foundational concepts of engineering-based security and the experiment underway to address the institutional and cultural problems previously described.

SECURITY FUNCTIONALITY AND ASSURANCE

There are two equally important aspects of protecting systems from adversarial and non-adversarial threats: security functionality and security assurance. Security functionality defines the safeguards and countermeasures needed to protect the organization's missions and the systems that support those missions. Security assurance is the grounds for justified confidence that a claim or set of claims about the systems has been or will be achieved (ISO/IEC/IEEE 2019). Justified confidence is derived from objective evidence that reduces uncertainty to an acceptable level and, in doing so, reduces the associated risk. Evidence is produced by engineering verification and validation methods. The evidence must be relevant, accurate, credible, and of sufficient quantity to enable reasoned conclusions and consensus among subject-matter experts that the claims are satisfied. Assurance is a complex and multi-dimensional property of the system that builds over time. Assurance must be planned, established, and maintained throughout the system life cycle (Ross, Winstead, and McEvilly 2022).

The determination of adequate security should be based on the level of confidence in the ability of the system to protect itself against all forms of adversity—that is, conditions that can cause a loss of assets. These conditions include threats, vulnerabilities, hazards, disruptions, and exposures. Adequate security cannot be based solely on individual efforts, such as performing functional testing, demonstrating compliance, or conducting penetration tests. Judgments of adequate security include what the system cannot do, will not do, or cannot be forced to do. These judgments of non-behavior must be grounded in sufficient confidence in the system's ability to correctly deliver its intended function in the presence and absence of adversity and to do so when used in accordance with its design intent. The basis for such judgments derives from well-formed and comprehensive evidence-producing activities that address the requirements, design, properties, capabilities, vulnerabilities, and effectiveness of security functions. These activities include a combination of demonstration, inspection, analysis, testing, and other methods to produce the needed evidence. The evidence acquired from these activities informs reasoning by qualified subject-matter experts who would interpret that evidence to substantiate assurance claims made. Assurance that also considers other emergent properties that the system may possess such as resilience to faults or adversarial incursions.




Figure 1. Balancing system cost, schedule, and performance with security

FOUNDATIONAL CONCEPTS

Systems engineering provides a foundation for a disciplined and structured approach to building assured, trustworthy secure systems. Security is an emergent property of an engineered system similar to safety, reliability, and resilience. As a systems engineering subdiscipline, systems security engineering addresses security-relevant considerations intended to produce secure outcomes. The engineering efforts are conducted at the appropriate level of fidelity and rigor needed to achieve trustworthiness and assurance objectives.

In security systems engineering for space systems, mission protection needs guide and inform the selection of security requirements and specifications (i.e., security functionality and assurance requirements). The protection needs focus on: (1) reducing the uncertainty associated with the space system’s capability (i.e., system behavior), and (2) controlling (i.e., reducing or limiting) asset loss due to adverse consequences. Adequate security involves a multitude of trade space and risk-based decisions that result in systems that are “as secure as reasonably practicable (ASARP).” Figure 1 illustrates the concept of balancing system cost, schedule, and performance requirements with protection needs.

The foundation of trustworthy, secure systems lies in the security design principles that are applied during the life cycle-based systems engineering process. The principles are described in NIST SP 800-160, Vol. 1 (Ross, Winstead, and McEvilly 2022) and include least privilege, least persistence, least functionality, defense in depth, reduced complexity, anomaly detection, mediated access, domain separation, and

least sharing. Security design principles are supported by system and cyber resiliency techniques and approaches as described in NIST SP 800-160, Vol. 2 (Ross, et al. 2021). The techniques and approaches are derived from the security design principles and include, for example, contextual awareness, adaptive response, coordinated protection, analytic monitoring, non-persistence, and monitoring and damage assessment.

NASA/JPL SUNRISE PROJECT OVERVIEW

In the prior section, it was articulated that the foundation of trustworthy, secure systems lies in the application of security design principles during the systems engineering life-cycle process of a system or mission. The expectation is that by doing so, the built system would exhibit improved resilience to faults and adversarial incursions. NIST SP 800-160 describes the security design principles, but not how those principles could be incorporated into well-established, well-exercised, systems engineering processes that underpin many operational systems and projects today.

Several questions arise when considering the application of the design principles such as:

- Where in the system life cycle should key engineering or trade space decisions be made for each security design principle (e.g., it may not be possible to apply certain design principles until critical system components have been built in the later phases of the systems engineering life cycle)?
- What approach or framework can systems engineers use to reason between operational resilience, safety, and security?

- What if the cost for engineering resilience into a mission turns out to be prohibitively high despite producing a quantum of resilience to an attack?

Furthermore, for operational systems like many of those in the U.S. critical infrastructure, the issues of cost, schedule and performance must also be part of the systems engineering decision parameters.

To explore this query toward achieving the desirable outcome of a more secure, resilient system, NASA/JPL in collaboration with NIST undertook a pilot experiment aimed at studying how the design principles for building trustworthy secure systems in NIST SP 800-160 could be incorporated into a well-established systems engineering process for space flight missions. The fundamental questions of interest for the experiment included:

- Can the security design principles in NIST SP 800-160 be integrated into the systems engineering life cycle of an operational system to produce a trustworthy secure system?
- How much improvement can be expected with respect to security when compared to the current approach that uses the NIST RMF and baseline security controls selected from NIST SP 800-53 (JTF 2020a) and NIST SP 800-53B (JTF 2020b)?

The mission selected for the NASA/JPL pilot was SunRISE, a composition of six CubeSats that work together to study solar activity. The science objective of the mission is to better understand how the Sun generates solar particle storms that can be hazardous to spacecraft and astronauts.

In undertaking the SunRISE systems security engineering pilot, a few fundamental challenges were identified in advance. Among the more prominent—the challenge of decomposing the design principles in NIST SP 800-160 into executable engineering actions that will integrate into the well-established, systems engineering life cycle of the SunRISE space flight project. Although the principles in NIST SP 800-160 have been established for some time, the constructs, models, processes, and frameworks needed to translate the principles into concrete engineering activities are largely absent in the literature, industry standards, and/or widely-accepted best practices.

Another notable challenge was that the SunRISE satellite project had to account for the pragmatic considerations of an engineered system deployed in a real-world context—namely cost, schedule and performance. In applying the NIST SP 800-160 design principles to an operational system, real-world constraints also had to be

considered in concert with the security and resilience of a built system. An operational system will necessarily include mission critical requirements, mission objectives, safety and reliability constraints, and other key considerations. All of these elements are necessary to achieve “mission resilience,” an emergent property of an engineered system similar to security. Consequently, the SunRISE pilot experiment not only measured the security properties of the engineered system, but also other mission-essential considerations such as cost, schedule and performance.

The next sections will describe the approach taken in the design of the SunRISE experiment where the intention is to illuminate and support the central premise of this paper: that a shift toward the addition of sound security systems engineering is needed to produce trustworthy secure systems that can more effectively address today’s adversaries. Additional technical details regarding the design of the SunRISE experiment including the significant number of engineering decisions and parameters that were employed, will be provided in future publications.

EXPERIMENTAL APPROACH

The following sections describe the experimental approach for the NASA/JPL SunRISE pilot project. These include the hypotheses, objectives, scoping criteria, and high-level methodology for the experiment.

Experiment Hypotheses

The overarching hypotheses for the SunRISE pilot established the basis for the experiment.

Hypothesis 1: Systems Resilience

- A systems engineering approach based on the application of the security design principles in NIST SP 800-160 produces a system that is more resilient and secure than a system that uses the traditional NIST RMF and pre-selected baseline security controls.

Hypothesis 2: Support for Risk-Based Decisions (Authorizations to Operate)

- A systems engineering approach based on the application of the security design principles in NIST SP 800-160 provides the necessary and sufficient assurance evidence to support credible risk-based decision making and the requirements for a system authorization to operate (ATO).

Hypothesis 3: Resources Required

- A systems engineering approach based on the application of the security design principles in NIST SP 800-160 can significantly reduce the level of effort, cost,

time, and resources required to achieve an ATO.

Experiment Objectives

The following objectives for SunRISE pilot are intended to test the experiment hypotheses:

- Demonstrate a working use case of applying the security design principles in NIST SP 800-160 to an actual flight project.
- Identify potential protection gaps in traditional cybersecurity approaches versus engineering-based security approaches.
- Identify potential security-related system design and implementation changes.
- Document the cost and effectiveness of engineering-based security.

Experiment Scoping Criteria

The experiment focused on the Ground Data Systems (GDS) component of SunRISE satellite system. The GDS is responsible for collecting and distributing the most valuable asset of the mission: the data. Several factors contributed to the choice of the GDS, the most prominent being that the SunRISE GDS operated in the cloud and could be easily replicated (i.e., creating a digital twin) for the purposes of this experiment.

Experiment Methodology

The high-level methodology for the SunRISE experiment is as follows:

- Identify the system under investigation
 - Identify a NASA/JPL mission that had already achieved its ATO
 - Identify a critical component (sub-system) of the mission
 - **Note:** The critical component selected needs to lie within the resource capacity allocated to the pilot (affordability) – the ground data system (GDS) for SunRISE
- Generate the replica of the system under investigation
 - Produce an exact replica of the SunRISE GDS (digital twin)
 - Twin A** – the original SunRISE GDS
 - Twin B** – the exact replica of the SunRISE GDS
- Establish the metrics
 - System performance (e.g., CPU resources, memory requirements)
 - Security performance (e.g., mean time to detection, mean time to remediation)
 - Programmatic (e.g., cost, schedule allowances, additional procurements)

1) Establish the baseline

- Conduct a functional evaluation of

Twin A and Twin B to ensure that the GDS functionality, resource usage, and system behaviors are identical between both instances (no attacks)

- 2) Select applicable security design principles from NIST SP 800-160
 - Principles selected based on SunRISE GDS architecture, mission requirements, and NIST guidance
 - Included Least Privilege, Least Sharing, Least Functionality, Mediated Access, Least Persistence, Anomaly Detection, Reduced Complexity, Defense in Depth
 - Also included resiliency techniques and approaches mapped to the security design principles
- 3) Design attacks against Twin B where the security design principles have been applied
 - The attacks were selected based on common security concerns for the GDS such as data exfiltration or the malicious modification of critical data
- 4) Design and implement defenses for Twin B
 - The security design principles from NIST SP 800-160 were used to design and implement the defenses into Twin B
- 5) Verification of GDS functionality
 - A functional evaluation was conducted on Twin B to verify that core the GDS functionality remained intact after the applying the NIST SP 800-160 security design principles
- 6) Execute attacks
 - Both Twin A and Twin B were subjected to the set of designed attacks
- 7) Collect and analyze results
 - Measurements for the selected metrics were obtained and the results analyzed.

The choice of a NASA/JPL mission that had already obtained its ATO was prompted by the need to compare the difference in security capability between Twin A (evaluated against NIST RMF and the SP 800-53 controls) and Twin B (integrated with defenses guided by the security design principles from NIST SP 800-160). The NIST SP 800-53 control evaluation for Twin A occurred during the mission’s Operational Readiness Review (i.e., toward the end of the mission’s design and implementation life cycle before launch). This means that Twin B did not “build on” a system already secured by the NIST SP 800-53 control evaluation to show improved security. Rather, the experiment is based on a Twin A and Twin B that were the identical standard NASA/JPL GDS design. The difference being that




Figure 2. Life cycle phases with Twin A and Twin B

Twin A moved forward to complete the NIST RMF/SP 800-53 controls assessment, while Twin B moved forward to integrate defenses based on the NIST SP 800-160 principles. Figure 2 illustrates the NASA life cycle process with the SunRISE GDS Twin A and Twin B.

Another notable point about the experiment was that the team that designed the defenses using the NIST SP 800-160 security design principles and the team that designed the attacks were separated from each other and did not communicate. This separation helped to ensure that the attacks on the SunRISE GDS were produced independent of the security design principles and implemented defenses.

It was also important that Twin B retained its native function as a GDS despite the modifications introduced by the NIST SP 800-160 security design principles. Consequently, additional tests were executed to continue the comparison of function and resource use between Twin A and Twin B. This was done primarily to ensure that Twin B continued to meet SunRISE mission objectives.

The final step of the process involved a detailed comparative analysis of data collected from both the original (Twin A) and redesigned (Twin B) GDS subsystems, allowing meaningful conclusions to be drawn about the effectiveness of the security capability achieved by the application of the security design principles from NIST SP 800-160.

INITIAL EXPERIMENT RESULTS AND INSIGHTS

The NASA/JPL pilot had only recently completed, and consequently, the discussion in this section describes the preliminary results and insights recorded. These insights tend to revolve around the role of the mission engineers and how they effected the outcome of the project. The initial results also highlighted the difference between the cybersecurity and mission

engineering disciplines with respect to the effort to build more trustworthy secure and resilient space systems.

Initial Result #1

A systems engineering approach tightly integrates security functionality into more aspects of a mission, better clarifying the impact and contribution of security to mission objectives.

In the SunRISE system, the traditional security control assessment for the ATO occurred after the mission systems had been designed, implemented and tested. Given the nature of the current security control assessment, this may make sense because at the earlier phases of the system life cycle, there are no implemented systems to evaluate. Implementation of the IT/cyber substrate for the SunRISE GDS is typically conducted in Phase C of a life cycle that begins in Phase A (design) and ends in Phase E (operations).

However, having the cybersecurity assessment occur in the later phases of the system life cycle facilitated a separation, or “siloing” of cybersecurity from the main SunRISE mission. The result was that the cybersecurity engineers who were engaged in the traditional RMF process did not have a strong understanding of the SunRISE mission, its objectives, or the engineering trades and decisions that contributed to the already built system. This lack of understanding resulted in one of the most noteworthy complaints from the mission engineers—that is, the cybersecurity engineers could not articulate how adversarial actions posed a risk to mission objectives, system capability, or how the security controls selected and implemented constituted a measurable reduction in risk to mission success.

Furthermore, because the traditional RMF ATO process was applied in the later phases of the system life cycle, the mission engineers did not have a strong under-

standing of the how the implemented security controls and artifacts contributed to the mission objectives and system capabilities critical to mission success. The controls and artifacts were perceived as incidental to the already built system.

One of the advantages of the NIST SP 800-160 engineering approach observed during the pilot project is that it engaged the mission engineers early in the system life cycle, specifically at the design phase. It catalyzed engineering questions and considerations with respect to the security-based mission failure implications associated with a specific NIST SP 800-160 security design principle being addressed. For example, consider the design and placement of sensors within a mission system—specifically, sensors that enable the engineers to detect anomalous application behavior with the objective of detecting mission failure. Instead of increasing system complexity and risk by adding cyber-specific intrusion detection sensors into the system, the mission engineers modified the sensors already in use for mission purposes, updated the concomitant operational processes for those sensors, and redesigned the sensor placement to make them dual purpose (i.e. detect potential faults and/or potential adversarial activity). This engineering activity occurred during the design phase and was prompted by the security design principle of “Reduced Complexity.” The application of this design principle ensured that security considerations were tightly integrated into the foundational design of the system and resulted in two significant outcomes:

- Because mission engineers were engaged in the modification and placement of the sensors for both reliability and security purposes, they understood the role of the sensors—that is, what part of the system and mission environment these sensors were monitoring and what the output of the sensors would mean to specific mission objectives. This provided the mission engineers with a stronger grasp on how to diagnose a problem or anomaly.
- Because the sensors were developed within the scope of the mission’s systems engineering process, sensor operation, maintenance and contribution to the mission workflow was tightly integrated. The mission engineers knew what the output of the sensors meant, they knew how to process that output, report the findings and perhaps more importantly they understood the impact of sensor failure to the mission.

The outcomes described above could

not be attributed to the traditional RMF approach simply because the approach as executed today, does not address how to integrate security considerations into the early phases of the system life cycle (e.g., the design phase).

In addition, the two outcomes above also support the following insights:

- A trustworthy secure systems engineering approach works well because the process involves tightly integrating security functionality into other mission-critical areas and not only on the cyber-related infrastructure.
- A trustworthy secure systems engineering approach is critical to enable rapid detection of adversarial behavior and diagnosis of potential adverse impacts and consequence to mission objectives and capability.

The first point notes that the integration of security into mission systems does not begin and end with implementing security controls only into the system's technical assets. The integration of security functionality must also include elements such as the mission's operational workflows, operational processes like the mission's Anomaly Resolution Process (ARP), and human resources who are able to understand the functionality of the security capability within the context of the mission, cost, schedule, performance, and maintenance considerations. All these elements are naturally addressed in the systems engineering process that properly establishes the security capability in the mission system. This is a fundamental reason why in the sensor scenario above, Twin B (using the NIST SP 800-160 security design principles in a life cycle systems engineering process) was observed to be more effective at addressing an adversarial incursion than Twin A (using the traditional RMF approach).

The second point notes that because the systems engineering life cycle engaged the mission engineers with security considerations from the start, the security functionality was incorporated into the workflow and processes of both mission engineers and operators. This means that when an anomaly or incursion occurred in the SunRISE experiment, the necessary steps to identify, diagnose, and remediate the issue were already "built in" as nominal mission processes, and could be executed rapidly and effectively by the mission team.

Evidentiary support of the two insights is suggested in the preliminary results from the SunRISE pilot:

- The mean time to detect the data tampering attack injected into both Twin A and B was reduced from weeks in Twin A to minutes in Twin B.

- The detection of mission data destruction was reduced from weeks in Twin A to seconds in Twin B.
- A malware-based malicious data tampering incursion executed on both twins was not detected with the traditional RMF approach on Twin A, but was detected with the trustworthy secure systems engineering approach on Twin B.

Initial Result #2

The traditional RMF approach can impede mission resilience and/or success.

The RMF approach applied toward the latter phases of the system life cycle lacks alignment with a mission's objectives and its systems engineering life cycle, creating a disconnected process. This disconnect meant that the mission engineers did not fully understand the security components introduced to satisfy compliance requirements or how to incorporate that security functionality into the space mission's operational profile. This introduced risk to mission resilience and success.

An example of when the RMF approach became an impediment to mission resilience and mission success was observed on the SunRISE pilot. When an anomaly appeared in Twin A and Twin B (an anomaly that was caused by a specific attack introduced to both twins), the mission engineers associated with Twin A saw the anomaly but were confused about what it represented and how they were to address it. Because the security components were introduced into the system late in its life cycle, the mission engineers did not understand the output of those components (such as a SIEM) or the semantics of that output with respect to mission failure/success. In short, the mission engineers did not understand what the presence of security components did to mitigate the threats to the mission or how to interpret the output from those components. It was unclear to the mission engineers where to look, how to understand the security audit data, and what it meant to mission objectives and capability.

Mission operations typically require the rigorous treatment of unknown events and anomalies using a well-established ARP. In the ARP, mission engineers are compelled to address the causal mechanisms underlying a given anomalous event. It took significant effort for the mission engineers to diagnose the anomaly introduced into the SunRISE GDS by trying to understand what the outputs from the security component introduced into the system. This significantly impinged on the natural mission processes that had to take place and consequently risked mission success.

It was observed that although the application of discrete technical controls as-

sociated with the RMF served the cybersecurity compliance requirements, it was less clear that they contributed to the overall success of the mission. Furthermore, the SunRISE experiment revealed that where cybersecurity engineers saw the adversarial threat as the primary motivating function for protecting the mission, the mission engineers saw the adversarial threat as merely one of several significant threats that could impinge upon mission success. The other threats would include structural failures, man-made disasters, human errors and so forth. In short, *cyber resilience did not equate to mission and system resilience*. This discrepancy in underlying concepts served to effectively block the successful integration of security into mission systems.

Initial Result #3

Small, inexpensive modifications engineered into mission systems can result in significant gains in resilience against a cyber adversary.

Two examples discussed in this initial result illustrate the consequences of incorporating security early into the system life cycle to address mission objectives. The first focuses on the security design principle of "Mediated Access" that involved the application of anti-virus scans designed by mission engineers. Although the traditional RMF approach did check that anti-virus scans were executed, it couldn't exercise the necessary depth of knowledge to check that the scans were executed on key system components critical to the mission. The mission engineers had that deeper knowledge of not only where in the system to apply the scan but also when in the lifecycle would a scan pose the least risk to mission objectives. Consequently, when they were engaged, they proposed a design of placement and workflow for the scans that were not only more effective than the ones assessed by the compliance-based approach, but also more economical in terms of initial deployment and maintenance costs. The solution designed by the mission engineers was more effective and less costly because it was targeted and intentional. It explicitly addressed the mission's critical assets and critical system life cycle phases (e.g., phases where external project partners were scheduled to deposit data into the mission's critical repository, each deposit was scanned before incorporation into the repository). This effective modification was small and inexpensive when engineered during the design and implementation phases, but it would have been an expensive addition after a compliance-based assessment.

A second example concerns application monitoring, a consideration under the "Monitoring and Damage Assessment"

cyber resiliency approach. Engineering the necessary logging capability to capture application telemetry for identifying adversarial incursions can constitute about 2 to 3 lines of code during the design and implementation phases (e.g., to capture CPU resource usage, memory usage patterns, application communication profiles, etc.). However, if this logging capability were to be added after a compliance-based assessment of a system that had already been designed, integrated, and tested, the cost would be prohibitive, and the issue would be delegated to the list of risk-based decisions that the project must make.

Lessons Learned

The most prominent observation from the NASA/JPL pilot project is that the integration of security into mission systems does not begin and end with the system's technical components. For operational viability, the integration of security functionality must include those areas that are naturally addressed within a systems engineering approach such as the mission's operational workflows, operational processes like the ARP, trained operators who understand the performance and functionality of the security components within the context of the mission itself and the associated cost,

schedule, and performance objectives.

CONCLUSION

The traditional RMF approach to cybersecurity and the associated ATO process works extremely well on enterprise IT systems that use mostly commercial off-the-shelf products. However, for certain types of systems being developed for high-intensity, mission critical operations such as NASA space flight systems, DoD weapons systems, and other high-value assets in the U.S. critical infrastructure, a systems engineering approach is needed to help ensure that security is treated as an emerging property of a mission system and integrated into the system life cycle. NASA/JPL conducted an experiment on the SunRISE satellite space flight system to determine if applying the security design principles from NIST SP 800-160 as part of a disciplined and structured system life cycle process, could result in more effective protection for the space system. After executing the traditional cybersecurity RMF process and completing the control assessments necessary to achieve an ATO, a comparison was made to the same system (i.e., a digital twin) that used a carefully selected set of security design principles from NIST SP 800-160. The initial results

were extremely promising with respect to the engineered system that embodied the design principles. By applying the security design principles early in the system life cycle as part of an engineering process, the SunRISE mission engineers had increased visibility into the system architecture to facilitate better placement of the selected security safeguards and allowed those safeguards to be more effective against an adversarial threat. The mission engineers were also able to reduce the complexity of the SunRISE GDS which also contributed toward achieving a trustworthy secure system that was more resilient. The initial results from the experiment prompted NASA to move into the second phase of the experiment, selecting a more complex space flight system and exercising additional security design principles from NIST SP 800-160. The complete SunRISE GDS results will be published and made available at the future publication. ■

ACKNOWLEDGEMENTS:

The content has not been approved or adopted by NASA, JPL, or the California Institute of Technology. Any views and opinions expressed herein do not necessarily state or reflect those of NASA, JPL, or the California Institute of Technology.

REFERENCES

- Anon. 2014. Federal Information Security Modernization Act (P.L. 113-283). s.l.:s.n. Available at: <https://www.govinfo.gov/app/details/PLAW-113publ283>. [Accessed 3 May 2025]
- International Organization for Standardization/International Electrotechnical Commission/Institute of Electrical and Electronics Engineers. 2019. ISO/IEC/IEEE 15026:2019 Systems and Software Engineering – Systems and Software Assurance – Part 1: Concepts and Vocabulary. Geneva, Switzerland: International Organization for Standardization.
- Joint Task Force Transformation Initiative. 2022. Assessing Security and Privacy Controls in Information Systems and Organizations NIST Special Publication (SP) 800-53A, Rev. 5, Gaithersburg, US-MD: National Institute of Standards and Technology. Available at: <https://doi.org/10.6028/NIST.SP.800-53Ar5>. [Accessed 3 May 2025]
- Joint Task Force. 2018. Risk Management Framework for Information Systems and Organizations: A System Lifec Cycle Approach for Security and Privacy. NIST SP 800-37, Rev 2, Gaithersburg, US-MD: National Institute of Standards and Technology. Available at: <https://doi.org/10.6028/NIST.SP.800-37r2>. [Accessed 3 May 2025]
- Joint Task Force. 2020a. Security and Privacy Controls for Information Systems and Organizations. NIST Special Publication (SP) 800-53, Rev. 5., Gaithersburg, US-MD: National Institute of Standards and Technology. Available at: <https://doi.org/10.6028/NIST.SP.800-53r5>. [Accessed 3 May 2025]
- Joint Task Force. 2020b. Control Baselines for Systems and Organizations NIST SP 800-53B, Gaithersburg, US-MD: National Institute of Standards and Technology. Available at: <https://doi.org/10.6028/NIST.SP.800-53B>. [Accessed 3 May 2025]
- National Institute of Standards and Technology. 2004. Standards for Security Categorization of Federal Information and Information Systems. Federal Information Processing Standards Publication (FIPS) 199, Washington, US-DC: U.S. Department of Commerce. Available at: <https://doi.org/10.6028/NIST.FIPS.199>. [Accessed 3 May 2025]
- National Institute of Standards and Technology. 2006. Minimum Security Requirements for Federal Information and Information Systems. Federal Information Processing Standards Publication (FIPS) 200., Washington, US-DC: U.S. Department of Commerce. Available at: <https://doi.org/10.6028/NIST.FIPS.200>. [Accessed 3 May 2025]
- Neumann, P. G. 2004. Principled Assuredly Trustworthy Composable Architectures, Menlo Park, US-CA: SRI International. Available at: <http://www.csl.sri.com/users/neumann/chats4.pdf>. [Accessed 3 May 2025]
- Office of Management and Budget. 2016. Office of Management and Budget Circular A-130, Washington, US-DC: Office of Management and Budget. Available at: https://www.whitehouse.gov/wp-content/uploads/legacy_drupal_files/omb/circulars/A130/a130revised.pdf. [Accessed 3 May 2025]
- Ross, R. et al. 2021. Developing Cyber-Resilient Systems: A Systems Security Engineering Approach NIST SP 800-160 Vol. 2 Rev. 1, Gaithersburg, US-MD: National Institute of Standards and Technology. Available at: <https://doi.org/10.6028/NIST.SP.800-160v2r1>.
- Ross, R., M. Winstead, and M. McEvelley. 2022. Engineering Trustworthy Secure Systems NIST SP 800-160 Vol. 1 Rev. 1, Gaithersburg, US-MD: National Institute of Standards and Technology. Available at: <https://doi.org/10.6028/NIST.SP.800-160v1r1>. [Accessed 3 May 2025]

ABOUT THE AUTHORS

Ron Ross**

RONROSSECURE, LLC

Wildwood, FL, USA

ron@ronrossecure.com

** Former Fellow, National Institute of Standards and Technology

Kymie Tan

Jet Propulsion Laboratory

California Institute of Technology

Pasadena, CA, USA

kymie.tan@jpl.nasa.govSystems Security Working Group [continued from page 14](#)

REFERENCES

- Dove, R., K. Willet, T. McDermott, H. Dunlap, D. P. MacNamara, and C. Ocker. 2021. "Security in the Future of Systems Engineering (FuSE), a Roadmap of Foundational Concepts." Paper Presented at the 31st Annual INCOSE International Symposium, Virtual, 17-22 July.
- Dove R. 2022. "Setting Current Context for Security in the Future of Systems Engineering." *INSIGHT* 25 (2): 8-10.
- Dove, R., M. Winstead, H. Dunlap, M. Hause, A. Scalco, K. Willett, A. Williams, and B. Wilson. 2023. "Democratizing Systems Security." Paper Presented at the 33rd Annual INCOSE International Symposium, Honolulu, US-HI, 15-20 July.
- INCOSE. 2021. *Systems Engineering Vision 2035: Engineering Solutions for a Better World*.
- INCOSE. 2024. Guide to Security Needs and Requirements. INCOSE-TP-2024-146.
- INCOSE. 2025. *Systems Engineering Competency Framework* (2nd Edition).
- Monat, Jamie P., and Thomas F. Gannon. 2017. *Using Systems Thinking to Solve Real-World Problems*. College Publication.
- NIST Joint Task Force. 2020. NIST SP 80-53 Revision 5 Security and Privacy Controls for Information Systems and Organizations.
- Ross, R., M. McEvelley, and M. Winstead. 2022. NIST SP 800-160 Volume 1 Revision 1 Engineering Trustworthy Secure Systems, NIST.



Your next giant leap is online

Earn your Master's in Systems Engineering

Purdue University's online Master of Science in Systems Engineering offers a flexible, interdisciplinary curriculum for professionals looking to advance their expertise in complex system design, analysis, and optimization. Developed with Purdue's Systems Collaboratory, this program emphasizes leadership, technical communication, and cross-disciplinary problem-solving, allowing students to tailor their learning experience to career goals while gaining cutting-edge knowledge applicable to aerospace, manufacturing, and defense industries.

- Control Systems
- Engineering Economic Analysis
- Game Theory
- Human Factors
- Machine Learning
- Multidisciplinary Design Optimization
- Practical Systems Thinking
- Project Management
- Reliability Based Design

#2

BEST ONLINE MASTER'S
ENGINEERING PROGRAMS
U.S. NEWS & WORLD REPORT 2025

LEARN MORE



P
PURDUE
UNIVERSITY