FAA Responds to FISMA to Improve Protection against Cyber Threats

Marshall D. Abrams, Ph.D., The MITRE Corporation
James Daum, Federal Aviation Administration

Overview

There are many roads to the nirvana of protecting Information and Information Systems. The Federal Information Security Management Act (FISMA) provides a rich travel guide of how to get there from here (for many values of here and there). This paper highlights several points down the road in one such journey. "Being secure" is a grail which we quest, but can never achieve as long as there are cyber adversaries in the world. Like the Knights of the Round Table we tell our tales from the viewpoint of the forces for good.

Our tale is set in a post Medieval Castle dominant like era where defense tactics were changing from walled cities to roaming teams of scouts with responding cavalries. It describes how FISMA is manifest in policy and acquisition at the Federal Aviation Administration (FAA) and how the FAA uses it as a basis to achieve acceptable cyber security risk. We detail how the FAA is improving protection against cyber threats by tailoring revised FISMA guidance to align with the FAA's mission. We highlight common controls, which are available for inheritance across the FAA, in the belief that common controls can uniformly protect many agency IT systems and contain costs. Policies and procedures are being re-examined and potentially revised, in response to the shift from naïve risk avoidance to risk management. Increased integration of cyber security in acquisition process is also highlighted.

This paper does not describe current protections. Such a description could provide useful information and intelligence to *adversaries* about the FAA's protection strategy or exploitable vulnerabilities.

Every agency has a mission. The FAA's mission involves the **safe and efficient operation** of a very complex cyber-physical and procedural system, the National Airspace System (NAS). The NAS consists of the existing air-traffic aviation systems and the transformation of those systems with future Next-Generation (NextGen) technologies. It involves complex technology, lots of kinetic energy, many people and computers–some trying to control the system and others consuming the services the system provides. A lot of Information Technology (IT) is employed in the system, some general purpose but most special purpose. The mission is performed using many information systems, some tightly coupled and some loosely coupled. The NAS enterprise is ever changing, i.e., always in a state of flux with new systems and capabilities coming in and old ones going out. The cyber security environment is ever changing, making the journey continuous. Protections that were previously adequate can now be breached. We must be continuously alert for new threats, but cannot relax protections against old ones. All these characteristics, and others, make the NAS very difficult to secure effectively and efficiently.

Every agency has business functions; every agency uses IT to conduct its business. These are identified in the FAA as non-mission functions. The distinction between the FAA mission system and the non-mission systems is not intuitively obvious. The distinction is complicated by organizational structure, spheres of influence, budgets, mandates, ingrained cultures, and good-old confusion. The determination of mission systems is largely dependent on the function of the systems, but highly influenced by the organizational responsibility of the systems. The mission systems (NAS) and business systems (non-NAS) at the FAA come under separate management, which converges only at the top of the agency. The exact

balance of power (budget), decision authority, and responsibility can sometimes make it difficult to secure the FAA's mission uniformly and efficiently.

FISMA

FISMA has a major role in setting the context for protecting Information and Information Systems at every federal government agency, but every agency is different. Like the Arthurian knight-errant, the FAA has many stories to tell about its journey through the realm of FISMA. We add our story to set the context for what follows. We focus on the standards and guidelines developed by the National Institute of Standards and Technology (NIST) to facilitate implementation of the legislation[1] that requires federal agencies to develop and implement information security programs. NIST's security standards[2] and guidelines[3] represent a comprehensive reference library of security concepts that enables the FAA to create effective security solutions by applying its domain knowledge to protect its Information and Information Systems assets.

FISMA applies to the NAS because it employs IT owned and operated by the federal government or by contractors working on behalf of the federal government. Other IT that are part of the U.S. Aviation System are industry owned and operated (e.g., by airlines and by service providers) are not part of the NAS and are not covered by FISMA. They are critical infrastructure, covered by the President's Executive Order 13636, Improving Critical Infrastructure Cybersecurity[4] and NIST's Framework for Improving Critical Infrastructure Cybersecurity[5]. The relationships, interactions, and interfaces between public and private IT adds to the NAS cyber security challenge.

NIST Special Publication (SP) 800-53, *Security and Privacy Controls for Federal Information Systems and Organizations*[6], includes a comprehensive catalog of controls that cover (almost) every threat known to NIST at the time of publication. The reader unfamiliar with SP 800-53 may be interested in a summary published by NIST in February 2014[7]. It also contains three sets of controls for federal agencies to use as starting points in selecting the appropriate controls for protecting their IT systems. These starting points, called baselines, depend on the potential adverse impact (LOW, MODERATE, or HIGH) if the system is compromised or breached. The scope of the impact is much broader than the specific information or information system; it encompasses the agency's operations, assets, personnel, other organizations and individuals, and the nation. Moreover, the impact must focus on the FAA's mission and the role that the IT system has in performing that mission.

After selecting the security control baseline, the FAA tailors the controls by adding or removing controls as deemed appropriate (e.g., removing controls for wireless communication when none is employed), and setting specific parameters in controls that require them (e.g., time period for an inactive session to time-out and what happens as a result). This tailoring process to achieve mission success gives the FAA flexibility in selecting and adapting controls.

---

[1] http://www.gpo.gov/fdsys/pkg/PLAW-107publ347/pdf/PLAW-107publ347.pdf

[2] http://csrc.nist.gov/publications/PubsFIPS.html

[3] http://csrc.nist.gov/publications/PubsSPs.html

[4] http://www.whitehouse.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity

[5] http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214-final.pdf

[6] SP 800-53, Security and Privacy Controls for Federal Information Systems and Organizations, revision 4, April 2013. PDF (including updates as of 01-15-2014): http://dx.doi.org/10.6028/NIST.SP.800-53r4

[7] http://csrc.nist.gov/publications/nistpubs/800-53-rev4/sp800-53r4_summary.pdf

Tailoring SP 800-53 at the FAA

Each individual information system or sub-system is the responsibility of an Information System Owner (ISO). The ISO staffs are guided by various forms of policy such as Orders, requirements, handbooks, and standards. The terms "requirements," "controls," and "countermeasures" are used more-or-less synonymously, often dependent on the context. We will use them in that manner in this paper.

The NAS was developed with extensive redundancy and built in controls for system anomalies. Therefore, many of the NAS systems are categorized as MODERATE impact. This determination is primarily based on safety risk, but also includes the architecture, topology, and past performance of the NAS. Although some NAS sub-systems might appear to be of LOW impact when viewed out of context, their connectivity mandates their consideration as MODERATE. Administrative efficiency and simplification also influence this determination. The NAS system staff provides guidance, currently called "requirements," to the NAS ISOs. These requirements are the results of tailoring SP 800-53. Selection and implementation of security controls is controlled on a system by system basis by the ISOs, resulting in security controls that provide various levels of protection, consistency, efficiency, and effectiveness.

The non-NAS systems are diverse in technology, application, and management; leading to determination of the full range of impact–LOW, MODERATE, and HIGH. All FAA system acquisitions must comply with FISMA requirements. Non-NAS systems also include systems that support regulator certification and flight safety activities. Because of the FAA includes the regulator organization in the same agency as some of the regulated activities it oversees some organizational boundaries must remain. About 20 of these systems are identified as HIGH systems. These systems support certification and flight safety activities where the vulnerabilities may affect aircraft safety or manufactures proprietary or intellectual data.

NIST guidance significantly changed when SP 800-53 *r*evision 4 was published April 30, 2013. Revision 4 supersedes Revision 3, published August 2009, reflecting the evolving technology and threat space. Revision 4 also contains a new appendix of privacy controls, and related implementation guidance that reflect the public demand for privacy of personal information. The Department of Transportation has adopted the revision 3 with the implementation of the DOT Cybersecurity Compendium as policy. The FAA has also adopted the DOT Cybersecurity Compendium as policy and is progressing with plans for integrating revision 4 into its existing policies.

The FAA commissioned an analysis of the security controls defined in NIST SP 800-53 revision 4. The study was intended to assess the current compliance of existing controls and policies with revision 4 and identify opportunities where its adoption could gain efficiency and protection with common controls. This study is essential to protect the operation of the current NAS and the future NextGen as well as other FAA information systems (non-NAS systems). Analyzing the security controls defined in revision 4 is the starting point for the development of information system security requirements. This analysis will help the FAA determine if the latest security requirements and security controls are being applied, as appropriate, to assure the delivery of FAA mission services.

Common Controls

The concept of *common controls* has evolved with the revisions of SP 800-53. Common controls are intended to be used (inherited) by organization units (e.g., ISOs) other than the organization unit that is responsible for creating and operating the control, called the Common Control Provider. The FAA is studying revision 4 to determine how to maximize the implementation of common controls, in the belief that common controls can protect many or most agency systems and contain costs. An emergent part of

this study was mapping controls to FAA policies and procedures to identify opportunities for improvement. The controls were considered from an agency perspective in order to identify those controls that can be made common throughout the entire FAA.

Implementation of common controls results in a security capability that is *inheritable* by one or more organizational information systems. Security controls are deemed inheritable when the ISO responsible for a system or component can utilize implemented controls that are developed, implemented, assessed, authorized, and monitored by other entities. Responsibility is the key concept in inheritance. When a security control involves responsibility outside the chain of responsibility for the information system, that security control is inherited. Many of the controls needed to protect organizational information systems (e.g., security awareness training, incident response plans, physical security, rules of behavior) are excellent candidates for common control status. The implementation of Service-oriented Architecture (SOA) based capabilities with NextGen creates excellent opportunities for common controls. By centrally managing and documenting the development, implementation, assessment, authorization, and monitoring of common controls, uniformity can be achieved and security costs can be amortized across multiple information systems.

Security controls not designated as common controls are considered *system-specific* or *hybrid* controls. System-specific controls are the primary responsibility of the cognizant ISO. Security controls have a *hybrid* status when one part of the control is common and another part of the control is system-specific. For example, an organization may choose to implement the Incident Response Policy and Procedures security control (IR-1) as a hybrid control with the policy portion of the control designated as common and the procedures portion of the control designated as system-specific.

Partitioning security controls into common, hybrid, and system-specific controls can result in significant savings in implementation and assessment costs, as well as a more consistent application of security controls FAA-wide. Additionally, the adoption of the common and hybrid controls at an enterprise level reduces gaps and vulnerabilities that may exist when applied to individual systems. (e.g. Security patches are often not performed consistently when applied at a system level of responsibility.) While security control partitioning into common, hybrid, and system-specific controls is straightforward and intuitive conceptually, the actual application takes a significant amount of planning and coordination. At the information system level, determination of common, hybrid, or system-specific security controls follows the development of a tailored baseline. It is necessary to first determine what security capability is needed before the FAA can assign responsibility for how security controls are implemented, operated, and maintained.

The identification and selection of common controls is most effectively accomplished on an FAA-wide basis with the involvement of senior cybersecurity leadership. The senior leaders are also in the best position to select the common controls for each security control baseline and assign specific responsibilities for those controls.

The determination as to whether a security control is a common, hybrid, or system-specific is context-based. Security controls cannot be determined to be common, hybrid, or system-specific simply based on reviewing the language of the control. For example, a control may be system-specific for a particular information system, but at the same time that control could be a common control for another system, which would inherit the control from the first system. One indicator of whether a system-specific control may also be a common control for other information systems is to consider who or what depends on the functionality of that particular control. If a certain part of an information system or solution external to

the system boundary depends on the control, then that control may be a candidate for common control identification.

A control implementation is clearly hybrid when the common control provider provides some data to the consumer, perhaps as part of an interaction, and the consumer takes some local action based on that data. For example, in single-sign-on the identification and authentication server passes an authenticated identify and associated meta-data to the application system which applies system-specific criteria for fine-grained access control. Another clear hybrid control situation is when the consumer ISO or authorizing official (AO) determines that the protection offered by the common control is not sufficient and determines that additional system-specific controls will reduce the risk to a more acceptable level. Common controls may be offered as services. Merely paying for a common control as a service does not make that control hybrid; nor does negotiating a service modification.

Security plans for individual information systems identify which security controls required for those systems have been designated by organizations as common controls and which controls have been designated as system-specific or hybrid controls. ISOs are responsible for any system-specific implementation details associated with common controls. Senior information security officers for organizations coordinate with common control providers to ensure that the required controls are developed, implemented, assessed, and maintained for effectiveness.

Information System Security Policies and Procedures

Information System Security (ISS) policies and procedures are the mechanisms for communicating up and down the agency what should be done and by whom. Written policy and procedures support a consistent and effective security posture. ISS policies and procedures are informed by the ISO's experience. Policy and procedures have many names and formats (e.g., order, instruction, standard, handbook, checklist) and range from mandatory to advisory. Policy compliance is essential to good management. The agency study of revision 4 identified places where policy needed to be re-examined and potentially revised, as well as areas where SP 800-53 addresses a security concern for which there is no policy.

Part of the changing environment for protecting Information and Information Systems is a shift from naïve risk avoidance (which is impossible) to risk management. A plethora of private sector and government-wide security breaches has made it clear that we must accept that bad things will happen and start planning what to do about it. NIST SP 800-39, *Managing Information Security Risk– Organization, Mission, and Information System View*[8], published March 2011, provides guidance. Part of risk management is determining the level of risk the agency is willing to accept in performing the mission or business function (alternately, the impacts to the mission or business function that the agency is willing to accept). Identifying this risk tolerance is both new and difficult. It may be an emergent property of the policy- development process. The question "How much security is enough?" obliquely approaches determining risk tolerance.

SP 800-53 benefits from the collective experience of all federal agencies and communities of interest. It contains countermeasures for most threats and vulnerabilities known to its constituents. We say "most" rather than "all" because the environment is so dynamic. Some of these threats and vulnerabilities may be hypothetical and not yet detected in federal Information and Information Systems. Note that they might exist, but are undetected. Since there is not a one-to-one relationship among countermeasures,

[8] http://csrc.nist.gov/publications/nistpubs/800-39/SP800-39-final.pdf

threats, and vulnerabilities; the control statements in SP 800-53 do not explicitly identify threats and vulnerabilities.

NIST provides a rich set of guidance for protecting agency Information and Information Systems. This is not a simple problem. There is considerable guidance in SP 800-30, *Guide for Conducting Risk Assessments*, revision 1 published September 2012. Understanding the relationship among risk and risk aggregation, threat, likelihood of occurrence, countermeasure, vulnerabilities and predisposing conditions, impact (i.e., the magnitude of harm), risk evaluation and uncertainty, and applying these considerations to the agency requires a specialist. These specialists work with other subject matter experts and cognizant agency managers to develop agency policy for applying the SP 800-53 control catalog within the agency.

The revision 3 baselines were instantiated by the FAA in various forms of policy separately for mission systems, and non-mission systems. The FAA has been stimulated by revision 4 to update existing policy and write new policy. As mentioned above, there is an emphasis on writing policy and implementing controls that apply as widely as possible. Policy guides investments, implementations, and operations across the agency. Acquisition is the tangible instantiation of policy for investments and implementation of new and reworked systems.

Applying FISMA at the beginning of the life cycle faces the challenge of making ISS an integral part of the acquisition system, complementing the ISS policy expressed in other normative documents. Work is underway to refine acquisition artifacts that will inform even the earliest resource allocation decisions (e.g., determining a rough order of magnitude budget) on ISS risk. Accompanying guidance will help service organizations (i.e., acquisition initiators) develop their artifacts and identify security controls as requirements for acquisition.

Increased Integration of Cyber Security in Acquisition Process

Congress granted FAA exemptions in 1995 from Federal Acquisition Regulations (FAR) and directed it to develop a unique Acquisition Management System (AMS)[9] tailored to the needs of the FAA. The AMS establishes FAA policy and guidance for all aspects of the acquisition lifecycle from the determination of mission needs to the procurement and lifecycle management of products and services that satisfy those needs. The AMS applies to the activities associated with the analysis of agency needs, determination of requirements, analysis of investment alternatives, establishment of investment programs, allocation and expenditure of resources, procurement and deployment of needed products and services, the in-service management of fielded capability, and eventual disposal of obsolete products. Investment programs are sponsored, and funded as investment decisions by the FAA Joint Resources Council (JRC).

The FAA Acquisition Executive Board (AEB) commissioned a Security Working Group to study potential changes to the AMS to better integrate cyber security, complementing the ISS policy expressed in FAA Orders and other normative documents. The AEB is a cross-organizational body, chaired by the FAA Federal Acquisition Executive (FAE), which assists and supports the JRC.

The ultimate goal of the Security Working Group was to define, document, and implement common, hybrid, and system security controls early in the system lifecycle as part of the acquisition process. This change was needed to provide more focus on the need for security controls and the identification of those controls as requirements earlier than what previously done during the Security Certification and

[9] http://fast.faa.gov/

Accreditation process prior to implementation. The work group identified three necessary tasks to achieve this goal: (1) define acquisition artifacts to inform JRC decisions on ISS risk, and ease transition from early acquisition to deployment by building a set of initial system authorization documents; (2) develop AMS guidance to help service organizations (acquisition initiators) develop their acquisition ISS assessment artifacts and allocate artifacts review to ensure they meet new ISS AMS guidance; and (3) update the AMS cyber security sections to incorporate guidance provided in NIST SP 800-37[10] and SP 800-39[11].

Background

SP 800-39 provides groundwork for 3-tiered risk management approach that progresses top-down from organization to missions to information systems. Its goal is to ensure that strategic considerations drive investment and operational decisions with regard to managing risk to the organizational mission, organizational assets, individuals, other organizations (collaborating or partnering with federal agencies and contractors), and nation. It describes a life cycle-based process for managing information security risk, including integrating ISS into the System Development Life Cycle (SDLC) of organizational information systems. Risk management activities take place at every phase in the system development life cycle, with the outputs at each phase having an effect on subsequent phases.

Early integration of ISS requirements into the SDLC is the most cost-effective method for implementing the organizational risk management strategy. Incorporating risk management into the SDLC ensures that the risk management process is not isolated from the other management processes used in the SDLC. Risk management (including information security considerations) is also incorporated into program, planning, and budgeting activities to help ensure that appropriate resources are available when needed—thus facilitating the completion of FAA program and project milestones.

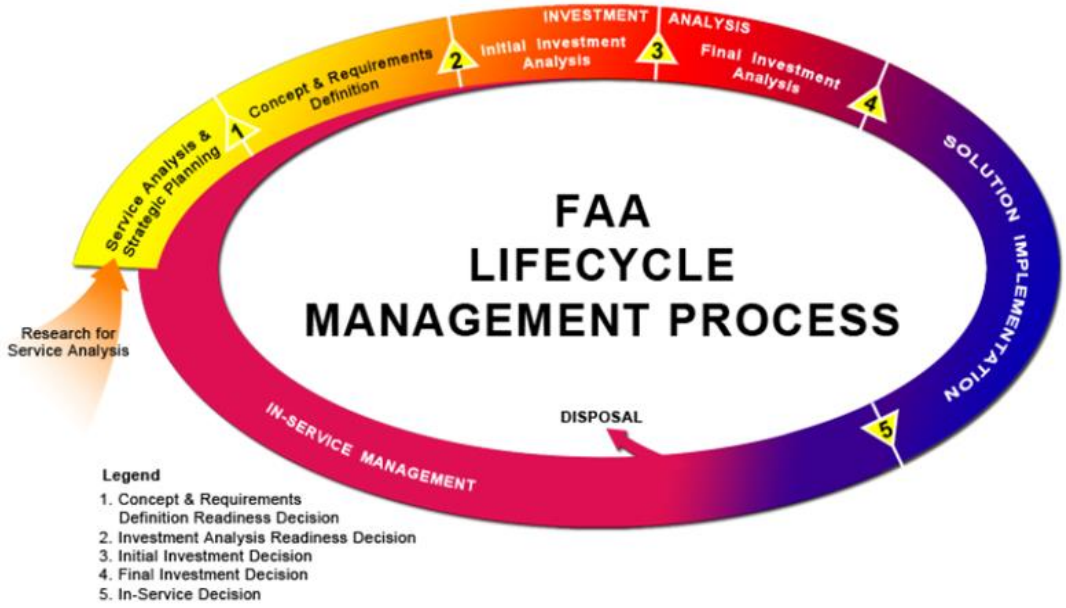

Figure 1. Acquisition Management System

[10] SP 800-37, *Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach*, revision 1, February 2010. http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-37r1.pdf

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

Incorporating ISS Requirements into the AMS

Lifecycle acquisition management is built into the AMS through a logical sequence of phases and decision points (see Figure 1). The FAA uses these phases and decision points to determine and prioritize its needs, make sound investment decisions, implement solutions efficiently, and manage services and assets over their lifecycle. The overarching goal is continuous improvement in the delivery of safe, secure, and efficient services over time. Application is flexible and may be tailored by the Acquisition Executive or Joint Resources Council.

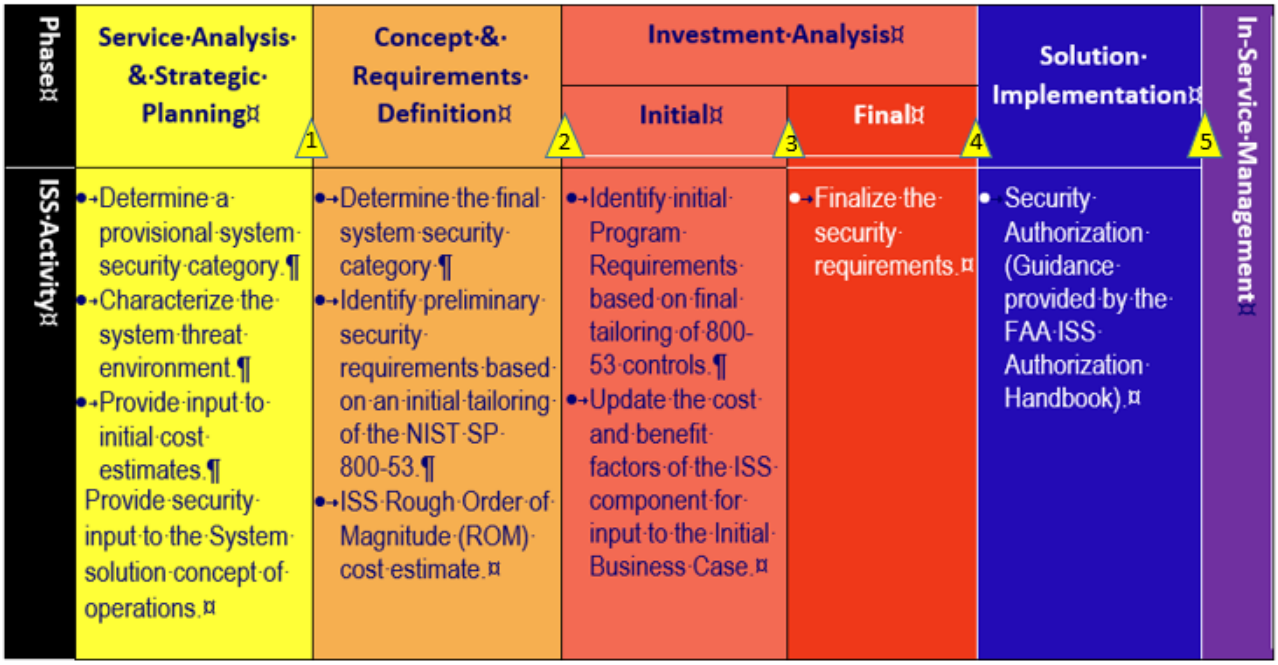

Figure 2. Alignment of ISS in AMS

Figure 2 shows the first four phases of the lifecycle process in a linear fashion with ISS activities for each phase, and the AMS decision gates between phases identified as numbered triangles (n). The four acquisition decisions are:

1. The Concepts and Requirements Definition Readiness Decision (CRDRD) by the FAA Enterprise Architecture Board (FEAB)
2. The Investment Analysis Readiness Decision (IARD) by the JRC
3. The Initial Investment Decision (IID) by the JRC
4. The Final Investment Decision (FID) by the JRC.

Requirements are developed incrementally by successive refinement so that the program office or service organizations initiating the acquisition are not overwhelmed with ISS requirements too early in the process or out of balance with the other requirements (NIST SP 800-53 ISS requirements are quite extensive). Some of the literature considers the ISS requirements as “assurance,” compared with the requirements of what the system must do, which are considered “functional.” Each AMS phase adds specificity to budget and schedule for solution implementation.

The AMS phases require a budget for the solution, which is also refined from phase to phase. The budget starts with an initial cost estimate, advances to a Rough Order of Magnitude (ROM) estimate and progresses to the specificity of an awarded contract. In the first phase, Service Analysis & Strategic Planning, a provisional system security category is determined, the system threat environment is

characterized, and security input is provided to the cloud suitability assessment and the System solution concept of operations. A team of ISS experts assists the service organization sponsoring the new system by applying previous experience with ISS acquisitions, technical knowledge, and understanding of the operation to provide an initial assessment of the level of effort that will be required for ISS.

In the second phase, Concept & Requirements Definition, the final system security category is determined and preliminary security requirements Identified based on an initial tailoring of the NIST SP 800-53. Additional guidance in the tailoring of security requirements is derived from National Airspace System (NAS) System Requirements[12] (NAS-SR), internal policy, and enterprise capability standards. (e.g. System Wide Information Management –SWIM). Often these are sources for identification or specification of common or hybrid controls. The initial identification of common, hybrid, and system controls in this phase is key to achieving the benefits described earlier. Experience and collaboration with ISS stakeholders and ISOs for common controls can help refine the acquisition budget and schedule.

In the third phase, Investment Analysis, the initial Program Requirements are identified based on final tailoring of 800-53 controls and application of agency directives and capability standards. The selection of common, hybrid, and system controls are evaluated and approved by ISS stakeholders. The cost and benefit factors of the ISS component are updated for input to the Initial Business Case. The final requirements, budget, and schedule are based on responses from a market capability survey, and experience with previously tailored FAA baselines. The third phase is increasingly focused on authorizing the implemented system for operation.

The fourth phase, Solution Implementation, follows existing policy and practice as stated in the ISS Authorization Handbook[13] and results in Security Certification and Accreditation.

Conclusions

This paper describes additional detail of several points of interest along the road that the FAA is traveling in its quest for a risk based ISS cybersecurity capability. The immediate journey takes the FAA down two avenues outlined in FISMA for all federal agencies to protect Information and Information Systems. These avenues are utilization of common controls and incorporating cyber security requirements in the earliest stages of the acquisition process.

Common controls are centrally managed to provide security control uniformity across multiple information systems, with the benefit of reducing security costs and seamless ISS protection. These common controls are inheritable by one or more information systems. The identification and selection of common controls is being conducted on an FAA-wide basis and specifically being applied to the acquisition process.

Early integration of ISS requirements into the SDLC is the most cost-effective method for implementing the organizational risk management strategy and ensures that the risk management process is not isolated from the other management processes so that budget and schedule decisions incorporate ISS

[12] https://nasea.faa.gov/file/get/2604

factors. Templates are being developed to define, organize, and standardize the information provided at the AMS decision gates.

In the spirit of our quest for security we should not be satisfied with the ISS changes to the AMS. These are great improvements that will make greater strides in identifying common controls and improving security of planned changes that originate out of an operational or efficiency requirement, but the effectiveness and benefit of common controls cannot be realized until an enterprise centric view is taken towards security. The FAA has begun to look at enterprise security concepts they may result in security based acquisitions and operations that will provide security functions to multiple NAS participants satisfying multiple security control requirements as a SOA capability. They are using the recently published NIST Cybersecurity Framework and using existing enterprise architecture products to help assess the system and functional security requirements of the NAS. The operational side is also using those products and its operational experience to identify areas where common controls will provide better security and increased efficiency.

The Cybersecurity Framework enables looking beyond FAA enterprise security concepts to improvements in the cybersecurity of the U.S. Aviation System. NIST published a companion Roadmap[14] to the Cybersecurity Framework that discusses next steps and identifies key areas of development, alignment, and collaboration. Both FISMA and the Cybersecurity Framework provide guidance and opportunities. While developed for federal agency use, FISMA standards and guidelines are frequently voluntarily used by non-federal organizations because of the flexible, risk-based, and cost-effective approach they offer. The Cybersecurity Framework and the NIST Risk Management Framework both seek to achieve the same objective – improved management of cybersecurity risk.

As with Jason and his Argonauts, Sir Lancelot and his Knights, and Dorothy and her band of unfamiliar friends, no great quest should be taken alone. The integration of systems, capabilities, and operations that is NextGen and the NAS knows no boundaries. The FAA and its private sector partners must become companions to complete this quest for improved security to achieve the FAAs goal of a safe and efficient NAS.

[14] http://www.nist.gov/cyberframework/upload/roadmap-021214.pdf