Beyond Compliance ---Addressing the Political, Cultural and Technical Dimensions of Applying the Risk Management Framework

Jennifer Fabius
Richard Graubart

Abstract

The Risk Management Framework (RMF) promulgated by the Joint Task Force provides organizations with a structured yet flexible approach to identify and prioritize the risks of depending on information, communications, and cyber-physical technologies; thus enhancing the ability to manage those risks. RMF implementation is in varying stages of maturity throughout the US Government. The RMF offers promise, but its implementation thus far raises questions and concerns about the direction the Federal government is taking to manage risk in a timely manner. Managing these cyber risks effectively requires organizations – and their mission or business elements, acquisition or procurement elements, and system owner-operators – to make political, cultural, and technical changes. This paper presents the benefits the RMF is designed to provide, challenges that organizations have faced, and recommendations to overcome those challenges and achieve the benefits.

Introduction

Security practitioners[1] use the term "risk management framework" (RMF) in multiple ways, depending on circumstances and the context of where it is being applied. Some use the term to refer to the collection of Department of Defense (DoD), Intelligence Community (IC), and Joint Task Force (JTF)[2] cyber security doctrine that provide a foundation for a common information security framework across the Federal government. Others use the term RMF to refer to the replacement for certification and accreditation (C&A) process. Some use it to refer to the six-step process shown in Figure 1 and described in NIST SP 800-37. Still others use it to refer to a shift in doctrine – the movement from a compliance approach to addressing security as a full lifecycle program to manage risk actively. Others use the term to refer to a combination of the above.

NIST describes the Risk Management Framework as a structured, yet flexible approach for managing the portion of risk resulting from the incorporation of information systems into the mission and business processes of the organization. The approach includes a six-step iterative process, as illustrated in Figure 1, informed by employing NIST, DoD, ODNI, and CNSS guidance which articulate risk management concepts and define specific process steps that organizations can tailor to meet their needs and constraints. The risk management concepts are intentionally broad-based with the specific details of

---

[1] By "security practitioners" we mean those engaged in applying any of the disciplines referred to as information security, information systems security, computer security, and cyber security, to systems engineering, business process engineering, strategic planning, program planning, or operations.

[2] The JTF refers to the collective effort of the DoD, Office of the Director of National Intelligence (ODNI), Committee on National Security Systems (CNSS) and the National Institute of Standards and Technology (NIST) to produce a core set of cyber security guidance documents that they all use.

assessing risk and employing appropriate risk mitigation strategies provided by the supporting NIST security standards and guidelines. Characteristics of the RMF, as noted on the NIST website, include the following:

- Promote the concept of near real-time risk management and ongoing information system authorization through the implementation of robust continuous monitoring processes;

- Encourage the use of automation and automated support tools to provide senior leaders the necessary information to take credible, risk-based decisions with regard to the organizational information systems supporting their core missions and business functions;

- Integrate information security more closely into the enterprise architecture and system development life cycle;

- Provide equal emphasis on the selection, implementation, assessment, and monitoring of security controls, and the authorization of information systems;

- Establish responsibility and accountability for security controls deployed within organizational information systems and inherited by those systems-- for instance, common controls; and

- Link risk management processes at the information system level to risk management processes at the organization-level through a risk executive function.
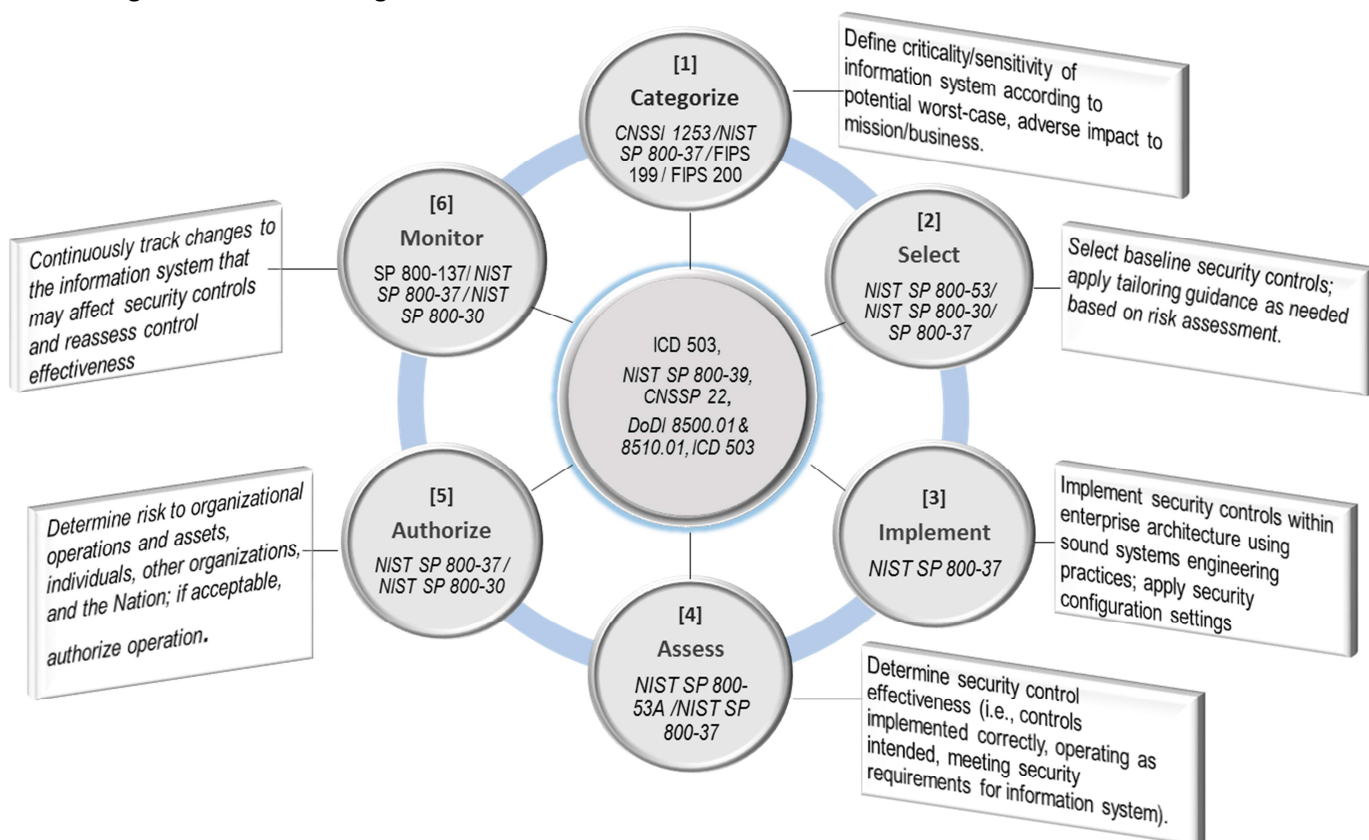

Figure 1: Risk Management Framework

The *intent* of RMF is to move away from a compliance-based approach to a risk-managed approach to cybersecurity. The drivers for this evolution include constrained budgets—including the need to "do more with less"—as well as the increasing prevalence of dynamic and sophisticated threats. Concurrently, the security community now recognizes what others have-- that security is not an end in and of itself, rather security enables an entity to fulfill its mission despite ongoing and successful attacks. Applying all possible security mitigations is cost prohibitive and interferes with the execution of missions (e.g., missiles on target, planes in the air).

RMF implementation is in varying stages of maturity throughout the US Government.   As leadership of Federal Departments and Agencies has changed, their commitment to implement the RMF as intended has fluctuated and the interpretation of what is needed has at times changed. Potential root causes for the issues noted in the paper stem from challenges associated with organizational culture, governance, and staff skill set along with the associated interplay between these factors.

Benefits

The RMF establishes common terminology for discussing cyber security risk across communities. The idea that security needs to be applied in a mission context provides a useful frame for discussions and decisions about managing risk. The RMF, as a structured, disciplined approach for assessing risk and determining appropriate mitigations to inform risk management decisions where security is a key but not the sole factor, promotes different kinds of discussions than those associated with a compliance model. Different components within the same organization or community often have limited perspectives on what is important and at times struggle to understand the validity and importance of other perspectives. Dialogue is an essential aspect of the RMF. With that dialogue, parties can better understand where they are in agreement and where their views diverge. Those additional insights allow for an authorizing official to make a more informed, risk based decision based on a richer set of information that historically afforded to that role.

Historically, each community (i.e., DoD, IC, civil) had its own set of security controls. For cross sector activities, work was required to find commonality between communities. As a result, hundreds of hours were spent conducting mappings between the security documents of the various communities. Furthermore, these mappings were not considered authoritative beyond the specific task that requested them; thus they were repeated for each new cross-sector project. The establishment of NIST SP 800-53 as joint defense, intelligence, and civil sector guidance ends the need for repeated mapping exercises by providing that needed and agreed to common set of security controls applicable across the Federal space.

Both NIST and CNSS have developed a series of *security baselines that provide a consistent foundation for selection of security controls*. These baselines articulate which NIST SP 800-53 controls should be selected by organizations. Baselines serve as a starting point and not as the minimum, and when used in that capacity, baselines provide a head start for identification of security needs. The concept of tailoring controls allows for addition, removal, clarification and where needed, modification of a control based on relevance to a system as well as organizational, technical, environmental, economic, and mission

priorities. Tailoring provides the flexibility to make the necessary adjustments after the selection of an initial baseline. This is a key element in making risk based decisions early in the lifecycle as encouraged in the RMF. Some entities appreciate the flexibility associated with tailoring; others express concern that it will lead to parties with similar interests diverging more than intended. To that end, CNSS established the concept of overlays as a way of providing a structured approach for tailoring when there are common technologies, information types, and/or mission settings. Tailoring and overlays reflect recognition by the security community that cyber security mitigations must be determined in the context they will be applied. The determination of the optimal security mitigations must take into account the various POET (political, operational, economic and technical) considerations that arise in selecting security control in a true risk-based process.

Successful execution of the RMF involves risk management activities throughout the system lifecycle, not an "accredit and forget it" mindset as has prevailed under previous approaches. Unlike previous government approaches for addressing cybersecurity risk, the RMF provides a legitimate avenue to accept the risk from addressing security needs differently than initially expected so long as it is done in a thoughtful manner. The reduced funding available to federal agencies reinforces the need for an informed risk-based approach such as what is promoted via the RMF. That *flexibility in a cost constrained environment* is key to managing what seem at times to be at competing expectations – increase security and use less funding. In this regard the RMF, the overlays, and the various doctrine that explicitly recognize tailoring supports leaders, as they engage with those who insist that all suggested security recommendations (e.g., the baselines in NIST SP 800-53 or the CNSSI 1253 baselines) must be followed. Ultimately the RMF is a vehicle to help leaders be more aware of the tradeoffs they will need to make because they cannot afford to address all possible security threats and still have sufficient funds available to support the core missions.

Challenges

In practice, *there are some large gaps between the RMF objectives and how organizations are implementing the RMF.* There is resistance towards viewing the RMF as an adaptable process. Resistance exists for a variety of reasons including but not limited to the following:

- Unfamiliarity about the flexibility inherent in the RMF,
- Limited engineering experience among many security practitioners and lack of familiarity with the concept of a trade space,
- Lack of supporting tools to help determine which safeguards are most appropriate,
- Pressure to remain within one's silo due to the political ramifications of convergence of security with other domains.

The maximum utility of the RMF cannot be achieved without overcoming such source(s) of resistance. For the purposes of this paper we are dividing the challenges into three categories: political, cultural, and technical.

Political

The vision for a unified information security framework was set in motion several years ago by the OMB, IC, DOD, and NIST. *Within the last few years a number of competing pressures have affected how the Federal government operates.* To cope, each community has become more focused on addressing the needs of their particular community. The continuing stress to do more with less within each community leads to questions as to whether a common vision remains.

Successful, efficient implementation requires solid governance as well as a culture that promotes communication, trust, thinking, and informed risk taking. Artificial limits in authority or willingness to trust peer organizations prevent organizations from being able to take advantage of the gains in efficiency that come from reuse and reciprocity.

Organizations need risk assessment, risk response options and risk-aware, mission-driven processes. Some of the implementation decisions made across civil and national security about applying the RMF are the greatest sources of perceived RMF problems. The RMF embeds risk assessments in each step – however the discussion of risk in most steps is so subtle that many do not recognize what risk-related activities need to occur. Across pockets of civil and national security community members, there are many who talk about and practice the RMF as if it were nothing more than security controls, security testing and evaluation, and continuous monitoring. While these concepts have a role in the RMF, in and of themselves they cannot and will not lead to risk management.

NIST designed the concept of *information security continuous monitoring (ISCM) to support risk aware, mission-driven processes*. Many organizations struggle with the following characteristics associated with ISCM:

- Defining specific roles and responsibilities, especially with outsourced services and providers,
- Determining what constitutes sufficient monitoring, and
- Evolving ISCM beyond compliance checks.

Monitoring for compliance is a factor but not intended as the primary reason for monitoring. Ongoing monitoring is a critical part of the risk management process. "In addition, an organization's overall security architecture and accompanying security program are monitored to ensure that organization-wide operations remains within an acceptable level of risk despite any changes that occur."[3]

Many officials fail to perform the necessary risk framing activities that inform the execution of RMF activities. Two of the most common reasons appear tend to be 1) risk aversion and 2) an unwillingness to articulate in writing their risk tolerance– the level and nature of the risk they are willing to accept. In particular, organizations with these challenges need to adopt a policy that assigns roles and responsibilities for framing, assessing, and managing cyber risks. Such policy should make explicit the relationships between these roles and responsibilities and those related to managing non-cyber risks. That is, cyber risk management must support enterprise risk management, which includes managing financial, operational (or mission), and existential risks.

---

[3] NIST SP 800-137 executive summary.

Despite the publication of NIST and multiple national security community documents over the last several years, many organizations believe that the transition towards the unified information security framework is being rushed. *The perception of a rushed transition exists because very little action was taken the first few years after the completion of the majority of the guidance.* Many organizations feel the pressure to transition to the RMF quickly because of the top-level guidance that previously was unclear or discounted. Many organizations feel constrained in their ability to meet the timeframes mandated. In some cases, transition timelines are expanded significantly because contract updates are not feasible in the immediate future and they refer to dated policies (e.g., DCID 6/3, DIACAP).

Cultural

Organizational change is a pre-requisite for evolving how the RMF should be implemented. *Organizational change is needed because many people fear the unknown*. With transition comes working with the unknown and to some that can be seen as too risky. Within the security community, adopting a risk *vice* compliance perspective is a significant cultural change. Where mindset is open, training accelerates adoption of the RMF. Where mindset is closed, the outcome of training has been mixed – with some people becoming more willing to embrace the way the RMF was intended to be used while others become more entrenched in their views that the process is cumbersome, bulky and ultimately a threat to security.

To date, multiple approaches for implementing the RMF have been tried by various organizations. While each has had differing experiences, one *common refrain heard is a need for "more" – more skilled staff, more resources, more time to transition and more training. The need for "more" has less to do with the RMF and more to do with the effects over time of underinvestment in staff capabilities*, as well as the inherent complexity of cyber security. The RMF roll-out brings to light some of these historical challenges and issues. Where an organization felt like they had staff who understood the engineering and operational aspects of their work, the transition experience was viewed as a net positive for the organization. Absent well-versed staff, there is a tendency to revert to compliance-oriented approaches for applying the RMF. Compliance is familiar and legal and policy doctrine provide support for compliance.

Legacy perceptions exist that risk is something that can be avoided, by taking actions that prevent the adversary from achieving a persistent presence within the organization. Such a view is unrealistic due to continually evolving adversary capabilities and intentions. *Some risks will always materialize, and they need to be managed*. Today's sophisticated adversaries are quite capable of achieving, and often expanding, a foothold in a system. Systems must be designed in such a way as to maximize their ability to achieve key mission functions, despite adversary presence.

A common belief permeating the culture within much of today's security and acquisition community is that if one spends sufficient time and energy up front addressing a cyber threat "properly" then little or no further action is needed. In reality, adversaries evolve and respond to defender actions rapidly and thus the interplay between adversary and defender has become much more dynamic than in years past. Therefore, effective risk management must be an ongoing process.

Technical

When users try to use baselines as the minimum and basis for compliance, challenges arise. The various security baselines often fail to articulate the operational or technical environment. Therefore *implementing baselines without tailoring sets up users for an unrealistic or unnecessary set of controls*. Many organizations have experience using the controls in NIST SP 800-53 and are learning actively from those experiences. Lessons learned about which controls are effective against different threat vectors and environments are beginning to be identified. Identification of which controls require greater organizational maturity, and/or more sophisticated set of defenders, etc. is beginning to be understood as well. However *most of this information is not captured in any common location* where practitioners can go to learn from those with experience. In addition many of the controls have various assumptions associated with them (e.g., assumes a physical infrastructure, assumes a high degree of persistence of data, or assumes that the organization is a government entity). But these assumptions are not articulated or captured in any knowledge base. This lack of a collective knowledge base and lack of automated tools that allow for meaningful mining and analysis of the knowledge base means that even if there were no political obstacles and the users have the appropriate risk management mindset they still *lack the information and tools that are needed to support making informed risk management decisions.* The situation is comparable to having trained and empowered medical professionals making decisions regarding prescribing medications/treatments to patients without any information and tools providing them information on the effectiveness, side effects or interactions of the various medications/treatments.

Most of the available *automated tools tend to be compliance focused tools*. For example, there are various continuous monitoring tools that support the monitoring step of the RMF; but they tend to focus on compliance and implementation status. Determining whether solutions implemented are compliant is an element of monitoring, but should not be the sole reason for monitoring. It is important to also monitor for other factors such as: 1) Effectiveness: are the mitigations (e.g., security controls) deployed effectively against the threats to which the system is exposed; 2) Relevance: are the mitigations relevant to the environment, have there been changes to the environment (e.g., new adversary TTPS) or the technology that impact the relevance of the mitigations.

In addition, the overall system security engineering process and system design and acquisition requires individuals making informed risk management trade-offs. Making these informed trade-offs requires two things.

- First, there needs to be a relatively authoritative body of knowledge with regards to the environment (e.g., which mitigations work well in a tactical environment). Understanding of the environment includes the following:
  - Nature of the likely threats (and associated adversary TTPs), and which are the most effective mitigations against these TTPs (and threats),
  - Relative cost of the mitigations, how effective are the mitigations, and
  - Operational considerations when implementing the mitigations.
- Second, there needs to be automated tools that allow system security engineer the ability to quickly mine such knowledge to facilitate, map them to the relevant NIST SP 800-53 security

controls and determine the dependencies among the controls, thus allowing them to make informed risk trade-offs.

Consider the changes in use of NIST SP 800-53. There are over 860 security controls in NIST SP 800-53, and in all probability this number will continue to grow in the future. Only a subset of the controls are applicable to any given system. The document is essentially a catalog of potential activities that one should do to implement a security program. However many people believe controls as-is are technical specifications. Others see the potential to use security controls as an input into the requirements management process. The number of controls and the relative merits and applicability of the controls is too much for any human being to keep in his/her head. Therefore, automated tools mining a well-maintained, shared database/repository, containing relevant metadata regarding the controls, are necessary to aid security practitioners in making informed decisions regarding the effectiveness, cost, and relevance of the various controls in different environments and different threat settings. Without these tools the security professionals are at a great disadvantage with trying to keep pace with the changing cyber threat environment and associated security mitigations.

Recommendations

The resource implications associated with RMF transition generate concern across multiple stakeholders. Resources are financial as well as human capital. With fiscal pressures come opportunities to think "smarter" and differently about an organization or community's approach to cybersecurity and risk management. With this recognition comes awareness that it is neither practical nor useful to employ a compliance approach to the selection of security controls. The selection of security controls needs to reflect the environment, the threat, and other operational and fiscal realities. The variety of systems is too diverse and the environments in which they are employed are too varied (e.g., in space vehicles, in mobile devices, in command and control environments) for a realistic, one-size-fits-all, solution.

Technical Recommendations

From a technical perspective what is required is a means of collecting and capturing an authoritative body of knowledge that can be reused by those making risk based decisions. To date solutions such as Collaborative Research into Threats (CRITS) only capture aspects of what is needed. The body of knowledge would capture information on various security mitigations (and associated security controls) and include information regarding which mitigations work best in specific environments, relative cost, maturity, and operational considerations, and information regarding which mitigations work best against which adversary threats and TTPs. Also needed, is automation that would allow security practitioners to mine, analyze and add to this knowledge base in analysis of possible mitigations (and associated security controls) in a timely manner. Some combination of databases and automated tools may also help decisions makers in determining their risk tolerance and risk thresholds. Establishment of risk tolerance and thresholds provide an organizational position that informs selection of risk responses.

The proposed knowledge bases and associated automation tools would improve the efficiency of controls selection and increase confidence of oversight authorities that the selected controls are appropriate and/or needed. They would also allow security professionals the ability to leverage the

experiences of others working in similar environments who have completed selecting and tailoring security controls. The proposed knowledge bases and associated automated tools would decrease the amount of time and effort put into security control selection.

Any tools, and any information generated by any of the proposed tools should be seen as decision-support for RM decision makers. Such information is used to maximize understanding of the options and subsequent risk based decisions made. The tools are **not** intended to be a different type of compliance vehicle – where users take the recommendations/results coming out of the tools and interpret them as gospel/mandates; that would simply replace the current flawed compliance approach for another flawed compliance approach.

There also should be a means to enable various organizations to share their experiences about the applicability/utility of the security mitigations (and security controls) and contribute to the knowledge base. In so doing this would expand and improve understanding.

Culture and Political Recommendations

Development of tools and databases is a relatively straightforward, although not simple solution. The greater challenge is changing organizational culture and politics to support risk management without retaliation for decisions that do not work out as intended.

The shift from a compliance mentality requires changing beliefs and attitudes that have been in place for over 20 years. Part of that shift entails recognition through to the highest levels of organizations that there is no single optimum security solution for all settings, and that any solution offered still has risks associated with it. The corollary to that is the realization that those decision-makers almost certainly will, over time, make incorrect decisions. This is not *carte blanche* to allow egregious or incompetent decisions. At the same time, one cannot expect individuals to make difficult decisions if they are looking over their shoulder constantly worrying about being second guessed.

One element needed to help achieve this change in mentality is more and better training for those responsible for making cyber security risk management decisions and for the security professionals who support them. Many people in these roles do not arrive at their jobs with the requisite training or experience to make informed risk based decisions. Addressing this shortcoming for those with the least experience may require comprehensive training that may entail a multi-month, if not multi-year training regimen[4]. For others it could entail training about the various elements of the RMF, risk assessment, underlying concepts embedded into the security controls, how to incorporate security needs into contracts, the nature of the advanced persistent threat, and/or practical exercises (e.g., tabletop and simulations).The nature of the training will vary depending upon the role of the individual along with their previous training and experience.

As noted above a key element of the transition from a compliance to a RMF approach is a willingness for decision-makers to acknowledge acceptance of risk. The RMF calls for the existence of a risk executive function that is responsible for making risk decision trade-offs. But not all organizations have to date

---

[4] By analogy, the training that a medic or EMT receives is different, and less extensive, than the training that a doctor receives.

implemented such a function. Moreover, such a risk executive function needs to be appropriately empowered to make the necessary risk management trade-off decisions.

The degree of risk that decisions makers are willing to accept will vary, based on the nature of the system, the mission which the system supports, default or actual organizational risk tolerance and decision-makers' beliefs about what is acceptable. Thus it is important to incorporate into the risk management process an explicit articulation by decision makers of the degree of risk that they are willing to tolerate. Having a clear articulation of the amount of risk that is acceptable and having that articulation captured will guide managers in making and framing their risk decisions.

Finally, as noted earlier in the paper, cyber security is not an end in of itself but rather a means to help achieve a mission. Therefore, we need to move from a perspective of cyber security being separate (sometimes opposing) activities that interfere with mission to one where cyber security is viewed as an integrated set of processes and activities that contribute to the execution of organizational and mission activities.

Conclusion

Applying the RMF as intended in the JTF publications is challenging work. It requires a fairly sophisticated set of skills, an appreciation for nuance, and an ability to operate in and navigate through different contexts. No one person does it all – it requires concerted effort by all the stakeholders to negotiate and operate across the organization, mission, engineering and operations and sustainment perspectives.

The RMF provides a structured, yet flexible approach for managing risk. When executed as intended, risk based decision-making at every step of the process allows for the options of acceptance, avoidance, transfer, sharing and mitigation of risk. Each type of risk based decision is valid—the decision of what to do with respect to risk should vary based on the diverse set of circumstances faced within a particular environment, organization, or community of interest. Effective risk management must be done in context of the strategic, operational and tactical imperatives facing an enterprise. If executed as a compliance vehicle or a separate silo, the benefits to be gained from the RMF will not be realized. Achieving the full benefits of the Risk Management Framework requires significant changes on the political, cultural and technical fronts.

References

Committee For National Security Systems Instruction 1253, *Security Categorization and Control Selection for National Security Systems*, May 2014.

Committee for National Security Systems Policy 22, *Policy on Information Assurance Risk Management for National Security Systems*, January 2012.Department of Defense Instruction 8510.01*, Risk Management Framework (RMF) for DoD Information Technology (IT)*, March 2014.

Department of Defense Instruction 8500.01, *Cybersecurity*, March 14, 2014.

Federal Information Processing Standards Publication 199, *Standards for the Security Categorization of Federal Information and Information Security*, February 2004.

Federal Information Processing Standards Publication 200, Minimum Security Requirements for Federal Information and Information Systems, March 2006.

Federal Information Security Management Act (P.L. 107-347, Title III), December 2002.

Intelligence Community Directive 503, *Intelligence Community Information Technology Systems Security Risk Management, Certification and Accreditation* September, 2008.

National Institute of Standards and Technology Special Publication 800-30, Revision 1, *Guide for Conducting Risk Assessments*, September 2012.

National Institute of Standards and Technology Special Publication 800-37, Revision 1, *Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach*, February 2010.

National Institute of Standards and Technology Special Publication 800-39, *Managing Information Security Risk: Organization, Mission, and Information System View*, March 2011.

National Institute of Standards and Technology Special Publication 800-53, Rev 4, *Security and Privacy Controls for Federal Information Systems and Organizations*, April 2013.National Institute of Standards and Technology Special Publication 800-53A, Revision 1, *Guide for Assessing the Security Controls in Federal Information Systems and Organizations: Building Effective Security Assessment Plans*, June 2010.

National Institute of Standards and Technology Special Publication 800-137, *Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations*, September 2011.