service design collective

Security & the Federal Risk Management Framework

Published July 28th, 2023: Version 1.2
Service Design Collective, Inc.
A public benefit company
Contact: rmf.project@servicedesigncollective.com

Table of Contents

Introduction

The Risk Management Framework is a set of security standards developed by the National Institute of Standards and Technology. It applies to all technical systems in the Federal government, except national security systems. The Risk Management Framework (the Framework or RMF) is also used by cities, states, and private sector companies. The Framework is a document outlining an approach to managing security. It is accompanied by several guides explaining how to apply the process to specific elements of technology.

The primary goal of this research was to understand both the theory and the practice of the Risk Management Framework in the Federal government. We wanted to understand what worked well and did not. We gathered stories of teams using the Risk Management Framework efficiently and effectively and asked participants about the future of technical security in government.

Service Design Collective gathered a team of five experts in the field of government technology with more than 50 years of combined experience. Our team read through the Risk Management Framework (RMF) and other supporting documentation related to software development and conducted two rounds of interviews.

Between August 2022 and June 2023, we spoke with more than twenty people, including Authorizing Officials, Senior Executives in security roles, policymakers, employees of private sector technology companies, and security consultants. Half of participants worked for the Federal government as senior executives or held the rank of GS-15, the highest level on the Federal Government's "General Schedule" pay scale. The rest worked at private technology companies or in Congress. Eleven participants held related positions in more than one role (executive branch, legislative branch, or private sector) in the last five years. All participants had worked directly on either the management of security policy or the implementation of the Risk Management Framework in the last year.

Private sector participants held titles that included Chief Executive Officer, Head of Compliance, Policy, or Transformation, Principal Software Engineer, Security Engineer, or Staff Engineer.

Government participants held Chief and Senior titles in roles that included Information Officer, Information Security Officer, Cybersecurity Engineer, Staff Member, Deputy Director, Information System Security Officer, Digital Service Expert, or Contracting Officer.

The Framework is a complex, personality-driven process. In theory, it provides a valuable foundation for security. In practice, it is unacceptably slow and expensive. It discourages modern security practices for all but the most inexperienced security professionals and delays or prevents the deployment of modern technologies that would help agencies achieve their missions. While some practitioners succeed in delivering effective security in an acceptable time and at a reasonable cost, that is not the norm. Even when the process is run efficiently, it produces sub-standard results.

The problem with the Framework is the process itself. If everyone makes rational decisions in their own self-interest, the Framework incentivises stakeholders to act in conflict with one another, rather than in concert. With changes, the process could be improved but in order to be successful long term, it must be radically simplified and rethought. Importantly, the biggest risk to relying on the Framework is not security failure, it is an inability to deliver critical public services and programs.

In this report we use quotes from NIST documentation as well as from our interviews. We chose to let our participants speak for themselves wherever possible under a promise that they remain anonymous when quoted. This is a work in progress and may be updated periodically to reflect new insights and additional research interviews. Participant quotes have been edited for clarity and anonymity. Changes are marked in brackets. Omissions are marked by ellipses.

The National Institute of Standards & Technology

The National Institute of Standards and Technology (NIST) has been responsible for cybersecurity and various iterations of risk management guidance for more than 50 years. Following the passage of the Federal Information Security Modernization Act of 2002, Congress tasked NIST with developing security standards and guidelines for all Federal systems (excluding classified systems). NIST maintains a thorough and approachable history of cybersecurity including information about the evolution of the Risk Management Framework that outlines their goals and strategies in technical security.

NIST has published a library of software risk management guides containing 24 core documents.[1] In addition, the Office of the Federal Chief Information Officer (OFCIO) publishes circular A-130, "Management of Information as a Strategic Resource." This collection of documents governs all aspects of Federal information security (see Appendix B).

The Framework is primarily focused on protecting data rather than securing systems themselves. This sets it apart from many other approaches to technical security, which see a system's functional availability as a primary objective. The Framework encourages systems to stop functioning rather than lose data (failing closed), whereas many modern software developers prioritize the continued availability of a product or service, even when a system is in distress or a breach has occurred (failing open). There are advantages and disadvantages to both approaches but this fundamental difference is the basis for significant tension when applying the Framework to more recent software systems.

[1] While some practitioners may argue that it is possible to navigate the authorization process with fewer publications (and some may reference additional documents), NIST has published these 24 documents specifically to support software security management.

Federal Agency Chief Information Officers (CIOs) create authorization procedures that govern the real-world implementation of the Framework. Authorizing Officials (AOs) in each agency use their agency's authorization policy and NIST publications to manage security, by issuing every system its own Authority to Operate (ATO). An Authorizing Official is responsible for understanding each document in the library, its role in the authorization process, and how to implement them. In theory, they should read and reference the documents for guidance and apply them consistently across their organization, and throughout the government.

The Risk Management Framework covers the basics

> *"Before RMF it was a crazier world. It got too crazy. So that's where the government kind of pushed back and said we need a standard way of viewing this."*

The Framework has helped many teams manage security in an organized, repeatable way and almost every practitioner we spoke with said that the Framework was useful at times. Some compared it favorably to a time before the Government Information Security Reform Act of 2000 (GISRA) when there was limited understanding around how to manage risk.

> *"I remember looking back with the passing of GISRA, we had authorized our first information system and it was foundational. It was really helpful because you have to think about what was there before it. There was no formal discipline corresponding to having understanding of your architecture, having an understanding in your control implementations or how things were done to specifically secure technical components, to operationalize management process oriented components. So there is value in what is there."*

Others noted that the RMF can prevent inexperienced technical teams from releasing insecure software.

> *"I have seen teams who have been blocked by the ATO process because they have no clue what they're doing."*

Other practitioners commented on the value of the Framework as a thought exercise or prompt to help them identify and mitigate issues they may not have considered.

> *"Best practices are time consuming ... but it also does force you to grow up."*

By far, the most common sentiment was that the Framework provides legitimate security benefits at a very high cost, both financially and in terms of labor hours. Many felt that the cost was too high.

> *"So it works from a security perspective. It's conservative and expensive, but it works."*

"Are there benefits to the security? Well, there are some, but those benefits aren't commensurate with the cost."

"It's expensive."

NIST writes for everyone but no one in particular

Looking objectively at each stakeholder in the Framework process: the developers and the private companies, the authorizing officials, the Information Security Officers, the technical writers, the policy writers, the legislators, and executives – all are following the incentive structure that the Framework creates. NIST writes the documents, includes different perspectives, incorporates new ideas, addresses new technologies, fosters public discussion, and creates documents for everyone to use and benefit from.

> Security and privacy control assessments are not about checklists, simple pass/fail results, or generating paperwork to pass inspections or audits.
>
> - Executive summary, NIST SP 800-53A

The unintended outcome of that public, inclusive Framework is that practitioners are put in a position where they feel they must use the Risk Management Framework as a checklist, despite the fact that NIST specifically says not to do so. In later documentation, NIST laments, but acknowledges, this common practice.

> An unintended and undesirable consequence of this has been that many security programs have focused on the individual controls as a compliance checklist, with little consideration given to how the controls work together to protect the confidentiality, integrity, and availability of information and systems.
>
> - 3.1.1 Supports Strong Systems Engineering of Security Capabilities, NISTIR 8011

If practitioners don't treat the Framework as a checklist, they may have to explain why they chose to include or exclude specific security overlays or security controls. This creates a situation in which, if a problem were to occur after a system was approved, it could be the fault of the Authorizing Official for not using a particular control on the list. Alternatively, If they complete the entire list of controls and then a problem arises, they can point to a completed checklist. In the latter situation, the blame is transferred from the Authorizing Official to whomever wrote the list or no one at all.

The Framework incentivizes using the complete set of controls[2] as a checklist, regardless of whether controls are relevant, or unnecessarily expensive to implement. Even when the outcome drives inflated costs, time delays, and even lowered security–no one is at fault. Everyone is acting applicably within the scope of the guidance.

[2] Outlined in NIST Special Publication 800-53

This may happen in part because NIST must write for the widest possible audience.

> *"NIST documents are for very broad audiences. Even if you just wrote it for Federal agencies, that would be broad enough. But they are often written for any organization that wants to use the documents. So, trying to be prescriptive with how to do things for any organization of any size, in any sector, in any country or state or province with customers from whatever parts of the world, with whatever types of sensitive data, whatever platforms they have, programming languages, and on, and on, it's impossible.... How will that get fixed in the future? Maybe through a lot of automation. But today there is no fix. Nobody has a solution to that."*

NIST carefully considered the broadest set of security standards and thoughtfully addressed them at length. But, by taking the widest possible approach to security, NIST has inadvertently subjected most practitioners to the highest effort approach to technical security.

> *"The way NIST created the 800-53 controls, they were trying to be as absolutely broad as possible and cover anything that could be considered an information system. And as a result a lot of the security controls sometimes...read as non-sequitur in the context of whatever system you're actually trying to get through the process."*

The Framework is growing larger & more complex

> There are over 860[3] security controls in NIST SP 800-53, and in all probability this number will continue to grow in the future.
>
> - Beyond Compliance: Addressing the Political, Cultural and Technical Dimensions of Applying the Risk Management Framework, The MITRE Corporation, 2014

Twenty years ago, following the Federal Information Security Modernization Act of 2002 and the development of SP 800-60, the Framework spanned approximately 400 pages, equivalent to the length of Herman Melville's Moby Dick. By 2008 it had surpassed Joyce's Ulysses. 10 years on, the Framework grew larger than Tolstoy's War and Peace. Today, the catalog of Risk Management Framework documentation[4] is lengthier than all seven books of the Harry Potter series.

Special Publication (SP) 800-53 alone is nearly 500 pages. NIST later released 800-53A, a guide to assess the controls in SP 800-53, that contains more than 700 additional pages. 800-53B, an addendum outlining the use of baselines for control selection, was later published in 2020 adding 85 new pages.

One example of how quickly the policy landscape can grow in complexity is the introduction of continuous monitoring. In 2011 The Office of Management and Budget (OMB) introduced continuous

[3] NIST SP 800-53Ar5 now contains 1189 baseline controls.
[4] Not including the recent Artificial Intelligence Risk Management Framework

monitoring but agencies struggled to implement it. NIST issued guidance for agencies in a new document, SP 800-137. In response to SP 800-137, OMB issued Memorandum M-14-03, giving three options for continuous monitoring. NIST then issued another document, NISTIR 8011 to manage those three options. One policy change led to three new policy documents.

> On January 26, 2023, NIST released the AI Risk Management Framework (AI RMF 1.0) along with a companion NIST AI RMF Playbook, AI RMF Explainer Video, an AI RMF Roadmap, AI RMF Crosswalk, and various Perspectives.
>
> - NIST.gov[online announcement], 2023

Most recently, to address the widespread adoption of AI, NIST released a separate, AI-specific RMF containing 72 new evaluation criteria, along with several supporting documents. In the last five years alone, the Framework has grown by nearly 1,000 pages.

> *"Writing more policy to have other orgs implement more policy to have other orgs implement their policy; that trickle down? I don't know. I'm skeptical that at the size of our bureaucracy it's going to make a change in a way that we need it to. Which is sad to say, and I want to be wrong. But it is my true answer that I think something has to give to where we break all the glass and we're just like, nope ... this doesn't work. Try again."*

Acknowledging the situation, a large agency CISO posed the question, "how has that 20 year old program ... matured? And to answer, it really hasn't. In fact, it's just grown. We're now onto NIST 800-53 revision five and if you go back through and look at the first version of NIST 800-53, the number of controls there have exponentially expanded, as have the baselines from low, moderate, and high."

The Framework struggles to keep up

The RMF struggles to keep pace with the demands of modern technology. The Framework's size and complexity make it increasingly difficult and time consuming for practitioners to manage.

> *"The challenge that we've seen over time is that the process hasn't fundamentally kept up with rapid change, particularly relative to how we manage environments, how we develop applications, and who developed applications."*

The catalog of NIST publications continues to grow larger and more complex as it attempts to keep pace with the speed of technical innovation. As technology is developed and deployed at an increasingly swift pace, a checklist-based management of that technology gets more complex. Complex regulation makes it harder, slower, and more expensive to authorize new software, compounding the problem. This has caused the gap between private sector and public sector software capabilities to grow dramatically.

> *"I don't even know how many different series of publications that NIST has. They have the SP 800 series, like 800-53, but there's also the 1800 series, the 1900 series. There's NISTIR, there's whitepapers, technical notes, and several other types of series."*

> *"The AI RMF is similar, you know, it is looking at risk management specifically in that AI context, meant to complement the Risk Management Framework and the privacy framework and other frameworks that NIST has."*

Many practitioners lamented the slow pace of the security assessments themselves. The Framework is a methodology that is used to create another document: an authority to operate (ATO). Estimates varied based on the complexity and risk analysis of the software in question, but it is common for ATOs to take six months, or longer, to receive. Several practitioners were involved in ATO approvals that took more than two years. Every practitioner we interviewed cited the need for the process to move more quickly.

> *"The ATO process is kind of impossible because ... it takes so long to write all those documents and get them approved and security just moves so quickly in terms of what the most secure posture is and, because of the bureaucracy, we just can never keep up."*

Even updates to existing systems can be laborious, leaving the government with outdated technology for extended periods of time. At best, system functionality may degrade slowly or useful features may remain unavailable. At worst, systems may be vulnerable to exploitation while solutions languish in the review process.

> *"The RMF does make it very, very difficult for people to keep things up to date."*

> *"I can find all the vulnerabilities, then it takes six months to fix it. And I'm like, what the heck is that? You totally defeat the purpose. You've created a process that takes forever to make a change."*

Every computer system operating in government is required to have an ATO. Put into perspective, the Department of Agriculture manages approximately 80,000 software programs. With a conservative estimate, where an ATO takes only three months to receive, the total time spent on managing security paperwork for the software in just one agency would approach 20,000 labor years. The cost of managing those same compliance processes quickly stretches into billions of dollars.

The Framework is well intended but not realistic

> *"The security is the security, the ATO process is a completely separate beast."*

NIST designed the Framework to be flexible and adaptable to different situations but practitioners we spoke to felt that flexibility is difficult to manage in practice.

> *"A lot of things that we do in the Federal government are so focused on compliance. And what we've been trying to do for quite a while now is, instead of focusing on compliance, to focus on the intent of compliance. Obviously there's an intention of the ATOs to make sure that people are building applications that are secure."*

Despite many attempts by our participants to implement the intent of the Framework, almost all practitioners ultimately used it as a compliance checklist.

Reinforcing this notion, when asked about the Risk Management Framework, practitioners primarily talked about Special Publication 800-53, the list of security controls. Few interview participants mentioned the Risk Management Framework itself (SP 800-37).

> Consistent with the flexibility allowed in applying the tasks in the RMF, organizations conduct initial control assessments during system development and implementation. Conducting such assessments in parallel with the development and implementation phases of the SDLC facilitates early identification of deficiencies and provides a cost-effective method for initiating corrective actions.
>
> - Implement, NIST SP 800-37
>
> Applying and assessing controls throughout the development process may be appropriate for iterative development processes .... This type of incremental assessment is appropriate if it is more efficient or cost-effective to do so.
>
> - Assess, NIST SP 800-37

NIST provides guidance for the management of security controls as part of agile or iterative software development but this message is buried in the literature. Even those familiar with it found it impractical to implement. Developers universally focused on software development and technical security as separate processes from applying the Framework. The Framework was considered a compliance exercise to be addressed after software development was complete.

> It is neither practical nor useful to employ a compliance approach to the selection of security controls.
>
> - Beyond Compliance: Addressing the Political, Cultural and Technical Dimensions of Applying the Risk Management Framework, The MITRE Corporation, 2014

Some Federal agencies did allow for Limited ATOs with fewer controls for minimum viable products. Even in the case of a Limited ATO, however, software was developed first and the Framework compliance process was appended at the end.

> *"It always seems like no matter how much I try to get something built in from the beginning, either scope creep or the customer changes directions or something happens and we always end up slapping it on at the end and I hate that."*

> *"I think most agencies, like I said earlier, they buy, they build then they bring in security. That's a huge problem"*

> *"I know ATOs don't solve anything, right? We're just going through compliance mechanisms"*

> *"You would get through seven, eight months of ATO and there would be no change in the security posture of your system. It was really more like security theater than anything else."*

Not only does NIST not require developers or agencies to use the Framework, they cannot enforce the controls. NIST is not a regulatory agency. While NIST lacks the authority to enforce standards, SP 800-53 is used by Inspectors General, the General Accountability Office, the Office of the Federal Chief Information Officer, and other oversight bodies to evaluate programs. Because of this, even though NIST is not a regulator, NIST guidance is made into de facto regulation in the eyes of Authorizing Officials, and every practitioner treats it as such.

> *"NIST is not a regulatory agency, so NIST has no power to require anybody to do anything."*

> *"You can make an informed risk decision [but] no one does that because the [Inspector General] comes in and says why didn't you do 800-53?"*

Ultimately, all interview participants felt that NIST was thoughtful and well-intentioned. However, there was an overwhelming sense that the Risk Management Framework was no longer an effective way to manage security.

> *"I think NIST is incredibly good at the academic, about the expertise around these things. They don't know how to turn what they do into a process or into something that people understand what to do with it."*

The Framework acts as an upper limit on security

> *"Unfortunately compliance has become so burdensome that sometimes it prevents updates to security because it takes so much time to audit it, so much time to prove it, that you don't have enough time to actually do the work itself."*

Interview participants noted that the Risk Management Framework failed to take into account how technologies are actually created.

> *"I literally looked at a requirement the other day that was about data spillage and it says all of this information about how to manage it on-prem and then it's like for cloud, maybe consider crypto erase, but that isn't really a viable solution for most customers."*

> *"I don't think NIST has really sat down and said, 'okay, this is what we care about. What are the ways we can meet that, that are compatible with the way development works now.'"*

And many felt that the Framework was a distraction to focusing on application development and more practical security concerns.

> *"There's going to still be a need for documentation and proof, but it can't be a 600 plus control, thousand page documentation because then you're spending all of your resources on creating words, not software."*

> *"You have literal droves of Federal employees that are hired specifically to create documentation, and you have zero Federal employees that are actual engineers. This is because of the perverse incentives that are created here."*

> *"One of the central things that's broken about the ATO process is that it doesn't consider the cost of not launching."*

While it is commonly understood that the Framework establishes a strong foundation for security, many experienced security professionals felt it also acted as a ceiling, limiting or discouraging the use of advanced security practices. Inexperienced practitioners benefit far more from the Framework, which provides a list of potential vulnerabilities and mitigations of which they may be unaware, than experienced professionals. Experienced security professionals felt that the control-based nature of the Framework limited critical thinking and that related FIPS standards, specifically FIPS 140-3 and FIPS 200, prevented them from implementing industry best practices.

> *"This control based approach does not promote critical thinking which is bad. Not just because the controls themselves become outdated and lag best practices and have bad outcomes on a micro scale, but on a macro scale it also leads to huge misses .... There's no control that is gonna have you thinking critically about what are the biggest risks facing my organization writ large."*

> *"Low, moderate, and high to a certain extent it's entrapment, right? It's essentially a race to the bottom. FISMA and NIST could adopt a threat-informed and more risk-based approach to security control selection that prioritizes the security of the individual information system and furthers broader agency and administration goals focusing on resiliency, shared services, and administration cyber priorities. Such an approach could be more responsive in aligning cybersecurity protection needs at system, agency, and governmentwide levels."*

"Yes, FIPS is a standard. Does anybody actually meet FIPS and be able to work in any kind of environment? Not really."

"Open SSH uses open SSL, but because it uses interfaces that are very old, we have to modernize it in order for it to use only the FIPs-compliant ones. Huge mess, huge undertaking."

Overall, the guidance discourages the use of creative, new, or potentially more secure technologies to manage risk.

Improvement may not be possible incrementally

All practitioners noted the need to update all or part of the Risk Management Framework, but many felt that incremental change was not possible. This led to a common theme that the RMF needs to be completely rethought.

"It's like turning the Titanic. It could happen but if you did ... the ripple effect would be very hard to predict and/or it would just monumentally fail."

"I don't see it getting any better absent congressional action to say this is the way we're gonna reinterpret agency risk management and cybersecurity technology."

"Most of our mission owners don't understand because they don't understand what the actual problem is, that these are just fictitiously invented bureaucratic problems. These are not real technical problems. These are just figments that we've put up in our way and I am of the belief that that is a political problem, period. And right now there is no incentive for our government to change that, absent a real threat."

One common understanding was that, when an emergency or a crisis required rapid response, the process was simplified dramatically or completely ignored in lieu of sound decision making. After, or in the absence of, a crisis, however, technical security reverted back to its less efficient state.

The Office of Management & Budget

The Office of Management and Budget (OMB) governs Federal security and privacy policies via the Office of the Federal Chief Information Officer (OFCIO) and the Office of Regulatory and Information Affairs (OIRA). OMB also governs the closely related fields of procurement and employee performance management via the Offices of Federal Procurement Policy (OFPP) and Performance and Personnel Management (OPPM). OFCIO plays the primary OMB role in technical security and publishes government-wide guidance in circular A-130, "Managing Information as a Strategic Asset."

A-130 instructs agencies to follow NIST publications, but also calls on agencies to "cost-effectively manage information security and privacy risks, which includes reducing such risks to an acceptable level." A footnote to Appendix A further states that "agencies must conduct tailoring activities in accordance with OMB policy." Like NIST, OMB encourages agencies to reduce the number of controls in the Risk Management Framework to "an acceptable level" and not to treat the Framework as a checklist.

A-130 also guides agencies to provide role-based training to security employees, but does not directly address competencies or training requirements for Authorizing Officials. Instead, it focuses on security awareness for the larger workforce.

OFCIO writes guidance but can not enforce it

> *"Does the Federal CIO have as much power to change agency behaviors as they think they do? I don't think so."*

In practice, OMB has little authority to implement their policies. They are dependent on agencies to develop authorization plans, select controls, and authorize systems. OMB even relies on agencies to provide data on their own performance.

> *"[OMB is] constantly operating in a blind spot when it comes to data because most of the data that they collect has to be self-reported from agencies regardless of if it's right."*

OMB's authority is most evident when the various offices within it coordinate to produce an outcome. For example, when the OFCIO defines the need for a new type of technical employee, such as a technical product manager, they are more likely to succeed in encouraging agencies to hire for the role when they coordinate with OPPM to define the role or designate senior executive positions. Likewise, they benefit from coordinating with the budget office to ensure funding is available to fill key positions.

> *"My biggest frustration when I was at OMB was that the people within the agency from the various management and budget side functions never actually worked together and coalesced around a clear set of objectives or priorities unless there was some chaotic thing that sort of pushed the agency to deal with it."*

Ultimately, the Office of Management and Budget has broad authority to advise agencies on how to manage technical risk and encourage them to improve government services but they lack the power to enforce their guidance.

> *"They're setting policy and telling people what to do. Does that mean people follow it? [Shrug.]"*

"Agencies have gotten very good at figuring out ways to create large budget pockets to manage what they're doing without OMB oversight. So they'll have large projects with 10 contracts under it, then they can maneuver that money and do what they want to. Rather than OMB saying we're gonna cut this contract by this amount because we don't want you doing that. Agencies have figured out the way to play that budget process as well."

"They have a lot of power because they're the White House. But it's sort of like the teacher in elementary school. You're gonna have your cliques that are gonna be doing stuff and you can influence them but there's only so much you can do to raise everyone up. It requires so many other factors. Getting those children to play well, to be nice, to have confidence .... I don't know how they ever will. That's part of the structure of our government, unfortunately."

"So I think part of it though is there is not enough money and educated people in IT and security in government. OMB puts out a new mandate every other week around new IT things that have to be done."

At the end of the day, agencies are in charge of security

"A lot of perverse incentives keep agencies from making smart, risk-based decisions that need not be there. And a lot of them are cultural, they're not legal, they're not even policy in a lot of standpoints. It's just worrying about the counterfactual versus taking the actions that an executive is charged with taking."

When asked, "Who is in charge of the ATO process?" respondents universally said that agencies ultimately made all security decisions. Despite NIST's mandate to create the RMF, and OFCIO's role in providing guidance, neither organization had a significant effect on how the Framework is used day-to-day. NIST cannot make agencies use the Framework the way it intended. OFCIO does not have the authority to make agencies tailor their controls to reasonable, cost-effective levels. Agencies will use the framework in a way that makes the most sense given their environment. This includes adapting to conditions and incentives unrelated to security.

Federal Agencies

Each Federal agency has a unique mission and programmatic goals separate from technical security. The Risk Management Framework process is often at odds with those goals. This conflict can take several forms. An agency may want to implement a new legislative mandate within a fixed timeline or respond quickly to an emergency situation by deploying new software on a timeline that the Framework does not permit. They may want to move away from legacy mainframe systems toward cloud infrastructure or digitize public services, in accordance with OMB policy, but find it easier to manage the Framework using legacy systems, servers, and mainframes.

In each instance, they will be confronted with a choice: move forward and contend with the ATO process, or remain on existing legacy systems that already have an approved ATO. They may choose to outsource a program or to simply go around the Risk Management Framework. These unapproved programs are referred to as “shadow IT” and are pervasive in government.

Even when the decision is made to implement new software, the ATO process often limits what that software can do; preventing useful features from being implemented or limiting the utility of service offerings. In other cases, the process ends at an impasse, with technical teams unable to provide adequate answers for controls that do not seem to apply to their software. This may result in a waiver or simply delaying the completion of controls to a later date via a Plan of Action and Milestones (POAM). Several interviewees remarked on systems with multiple unresolved POAMs that have remained unresolved through several rounds of authorization reviews over months or years. These are known issues that are simply never addressed.

The Framework incentivises legacy technologies and imperfect solutions

> *“So right, right now, the ATO process basically provides no incentive to move off of legacy systems because the mainframe has an ATO so we're just gonna keep it running forever. There's not really a prompt to reevaluate whether moving to a more modern system would actually improve the security of the system.”*

Programmatic decisions are often made for practical reasons, such as simplicity or ease, as opposed to technical capability. Because the Risk Management Framework process is lengthy and complicated, agencies sometimes choose to expedite the security process rather than try to fulfill all of the requirements of the program. In multiple interviews, we were told of times that a program had chosen to work with a previously approved technology or legacy system even though it only met a fraction of the technical needs. These decisions were made to avoid having to go through a new ATO approval process.

> *“Usually the government offering is like two to three steps at a minimum behind commercial [offerings].”*

> *“You can't really use the new features. So the code that's there doesn't benefit from the improvements we've made over the past 20 years.”*

The technology industry moves quickly. New features and capabilities are added to products over the course of weeks or months. The ATO process takes months to years. Over time, this disparity puts program managers in a difficult position – they must choose between using outdated tools or failing to deliver services.

This incentive to continue using existing systems has played a role in the continued use of critical Federal technology systems that are more than 40 years old. Reliance on such legacy systems is

both unwise and expensive. Approximately 80% of the Federal government's $100 billion IT budget is spent on maintenance for existing systems.

> *"Unfortunately compliance has become so burdensome that sometimes it prevents updates to security."*

More importantly, the reliance on certain legacy systems, such as those at the Center for Medicare and Medicaid Services and the Social Security Administration, can lead to more than just inconvenience, increased cost, and poor service delivery. They create nation-wide programmatic, financial, and social vulnerabilities. Were they to fail, agencies would be hard pressed to deliver any service at all, actively harming individuals, the economy at large, and the public's faith in government.

Agencies struggle to implement the Framework effectively

> *"The Risk Management Framework, I believe, in theory, is still fine and mostly valid as a way of thinking through a decision making process when it comes to IT and procurement and budget things like this. But no agency actually manages risks."*

Several factors affect an agency's ability to implement the Framework but even the most successful Chief Information Security Officers struggle to balance the financial and time costs with real-world security advantages. Agencies often lack the technical security talent to understand where they can derive value from the Framework and when to skip unrelated or unhelpful security controls.

> *"We could spend a year or two giving an ATO, we still would not achieve anything."*

Every agency security professional we interviewed felt the Framework was a compliance exercise. Some recognized value in the Framework, but struggled to balance the theory of the Risk Management Framework with the real implementation challenges. Most attempts to modernize it or integrate it into development were ultimately abandoned.

> *"I would make it more agile. I think it needs to fit how you build, I mean not the other way around. I mean when you're driven purely by compliance, that is not gonna fit the best practices for UI, UX, for engineering, for research, for ... nothing. It just enforces a different, and to me, outdated way of building which is; you write it all down and (essentially waterfall methodology) gather all your requirements. And you build out the entire system on paper, get it authorized, then you actually build it, right? So it's years before it actually touches a user and you realize, 'oh this didn't work, we need to change it.' Oh, you can't because that's what your ATO said."*

The Framework hurts productivity

> *"It goes without saying maybe, but I'll say that I think a more serious constraint is actually, at least in this and maybe other agencies as well: We can't use all of these software and tools in the workplace that are actually relevant to us understanding how to do our jobs better."*

NIST explicitly states that the Framework should not impede an agency's mission:

> Further, information systems process many types of information. Not all of these information types are likely to have the same security impact levels. The compromise of some information types will jeopardize system functionality and agency mission more than the compromise of other information types. System security impact levels must be assessed in the context of system mission and function as well as on the basis of the aggregate of the component information types.
>
> -SP 800-60, 4.3 Step 3: Review Provisional Impact Levels and Adjust/Finalize Information Type Impact Levels

In spite of that statement, there are CIOs that have delayed large programs, at great cost to the mission of the organization, in order to complete paperwork.

> *"It was like this tiny, tiny thing that was not a security issue that caused months of delay in the launch of this big system. Tens of millions of dollars invested in the development of the system that the CIO shop was just not willing to sign off on."*

There is no explicit instruction on how to prioritize mission over security. The Risk Management Framework also does not consider opportunity costs built into the speed of the Authority to Operate process or its effects on the way the Federal workforce uses technical tools. A year spent on ATO paperwork is a year in which slow, outdated, insecure systems may stay in operation. This delay becomes a cascading effect as the Framework gets larger and more complex.

When systems are authorized, oftentimes certain functions or features are disabled or, in the case of in house development, are never built. This can reduce the effectiveness of the solution and slow productivity. Other security requirements can slow or disable critical hardware, such as laptops.

> *"[The agency] was moving to Office 365... they have cloud based tools that you can collaborate in and all those things. But when Office 365 rolled out they turned off all of those features. What is the point of having Office365? It was like you're gonna have all of these features completely turned off."*

The public is also heavily impacted by the Framework process in their day-to-day interactions with the government.

"So if we don't get our house in order, if we don't remove all these points of friction, there's no way [my agency] or anybody else is going to deliver a pleasant experience to the beneficiaries or any of the other stakeholders."

Inconsistencies exist even within agencies

In attempting to outline every situation and technology that could be delivered in all of government, the Risk Management Framework creates inconsistent outcomes across the Federal government. Inconsistencies exist even between different offices inside a single agency. This has a direct effect on everything from system performance to the ability to collaborate within and across agencies. This is true regardless of whether or not software is off the shelf or custom built.

"If DHS is doing it, why is FEMA doing it differently if [FEMA is part of] DHS? It's another operational division. They own their process, it's their own people. I mean it's always, in my mind, politics and fiefdoms."

Most participants commented on the unpredictable length of time it takes to complete the Authority to Operate process. According to our interviewers, the average time it took to receive an initial ATO for systems categorized as moderate or high was over one year. On a few occasions, however, participants stated that they had received an ATO in days or weeks.

Receiving an ATO in a short period of time was more likely when a number of factors aligned. Combined, a strong sense of urgency, an Authorizing Official with knowledge of the system under review, and a willingness to accept uncertainty often led to quicker approvals. Because such alignments are rare, there is little consistency in how long it takes to receive approval.

"So to fix the situation, a high industrial control system, we put together all of the paperwork for that in about two weeks. The person doing the work on that was not ready for their assessment, that was happening in two weeks. So, we dropped everything and it was me and two other people. And this was probably easy because we already had background on the system.... So we were very, very familiar, more familiar than we wanted to be. So it was fairly easy for us to knock that out and, sure, we ended up with a lot of POA&Ms. There were a lot of them. We just didn't know the right answer when we were documenting it."

Some agencies have developed streamlined or "lightweight" ATO processes that improve both efficiency and consistency for certain types of systems. Some, for example, prioritize areas of focus and associated SP 800-53 controls in their agency authorization process so that developers and Authorizing Officials know where to focus their attention. For systems categorized low and moderate, some agencies have developed Lightweight Authority to Operate processes that authorize systems for a short period of time based on a reduced number of controls. At the end of the trial period, systems must still receive a traditional Authority to Operate, but the process provides developers an opportunity to put systems into production in a

way that is more closely aligned with modern software development. This lightweight approach also encourages teams to think about security during development rather than applying controls only at the end.

For commercial, cloud based Software as a Service that does not contain personally identifiable information, the FedRAMP program has a lightweight approval process called Li-SaaS. Li-SaaS applies a tailored set of controls and, more importantly, allows companies to attest to many of the controls rather than submitting detailed security paperwork making the process much faster.

In all of these cases, agencies are still constrained by the Risk Management Framework, including categorization and the application of specific security controls. Even so, agencies can encourage a more restrained use of the Risk Management Framework and communicate areas of concern that can speed up development.

As these examples show, however, as agencies strive to make the Authority to Operate process faster and more consistent, they become applicable to fewer technologies. Lightweight agency processes only apply to systems categorized as low and moderate. Li-SaaS applies only to a narrow band of applications that do not collect personally identifiable information. In addition, both processes still require a significant number of controls. Li-SaaS, for example, contains more than 200 controls and in all cases, Authorizing Officials can add additional controls as they see fit.

In some large agencies, even the agency CIO has limited control over how security decisions are made. They can narrow the scope of the authorization stage with agency-wide guidance, but must continue to rely on individual Authorizing Officials to manage the process.

Authorizing Officials

In practice, the Framework is managed by agency Authorizing Officials. Authorizing Officials are responsible for an overwhelming volume of work. Complicating matters further, security decisions made within the larger incentive structure of the Federal government are often driven by fear rather than security. AO's are often pressured from above and below; from development teams struggling to gain authority to operate to senior executives and political appointees trying to meet deadlines and budget goals.

Authorizing Officials work at the functional end of a policy process that is continuously evolving and changing. These changes come from new policies, technological advances, personnel turnover, process updates, and other factors that contribute to a fluid work environment. They are, more often than not, less technically skilled than the development teams they govern. The combination of a fluid implementation environment and an asymmetry in security experience causes significant problems in the practical implementation of the Framework.

AOs have a difficult job. They must categorize, review, and continuously monitor hundreds, if not thousands, of technical systems. They must take into account their agency's needs and mission, timelines, budget costs. They must monitor and reevaluate the broadest risk environment, from malicious software, to social engineering, and building security.

Authorizing Officials lack skills & need training

The most common refrain from interview participants was that Authorizing Officials do not have the basic technical and security skills to effectively understand threats, track technical improvements in the market, or manage risk.

> *"Most of the time the compliance folks just have no idea what they're talking about. They're getting asked to talk about tech they have no cognizance of, and then it just results in a combative relationship where neither side really wants to talk. It's like pulling teeth on all ends.*
>
> *"I've never met [an AO] that is actually technical."*
>
> *"If they don't have good judgment or and are not empowered to use that judgment, then it's not useful."*
>
> *"The guy who finally signed our ATO was not actually that competent."*
>
> *"The assessor had no idea. Lack of technical expertise."*
>
> *"The Federal government's approach has been, 'let's just take unqualified people and put them in those positions.'"*
>
> *"It would just work much better if you had people who had more technology experience or hands on security experience."*
>
> *"I have [application developers] responding to facility controls and it's like, 'What do you mean'? I don't control the building. This has nothing to do with me.'"*

Despite a common belief that most Authorizing Officials lack the skills necessary to be effective at managing technical risk, interviewees spoke sympathetically about insufficient training and formalized skill development for AOs in the Federal Government.

> *"Our [Information Security Officer] was a complete stickler because he didn't fully understand. He's just a guy who had a job in another state and probably did a year of training or something. He wasn't a cybersecurity expert and so he didn't have the confidence in really being able to tell what deserved a waiver and what didn't."*
>
> *"It speaks to a larger problem within the Federal government, which is how do we promote people, right? You take someone who's a really good budget examiner and you*

promote them to be a manager of other budget examiners, but they're not a good manager. They were just a really good budget examiner. And so, you know, then a lot of our SES of various agencies, which are typically the people that are getting assigned as an authorizing official, they're not necessarily what we would call in the private sector 'qualified executives,' they just know how to do their thing really, really well. And no one has ever sat down with them and taught them about managing risk at an organizational level. So, asking them to do that job is a little bit like asking them to fly an F-14, like they've never been trained how to do it. They can be the smartest person in the world. If no one's ever taught you how to do it, you're definitely gonna crash."

"A lot of those documents are written for lawyers. They're not written for practitioners. Like even 800-53; I realize that NIST is well-intentioned and wants these to be documents that folks can action. But no non-lawyer can take this amount of data and do a meaningful thing with it. That's not a skillset that the average person walking down the street has. [AOs] would have to say, 'let me take this hundred page document and make it real', right? That's literally a lawyer's skillset."

Authorizing Officials are experts in paperwork, not security

There is no official position description for an Authorizing Official. The Cybersecurity and Infrastructure Security Administration does post recommended traits on their recruiting website, but the job description of an Authorizing Official varies widely between agencies. Unlike roles with similar duties and responsibilities, such as Procurement Officers, there is no accreditation process to become an Authorizing Official. Training is sparse and inconsistent.

"There should be a certification process for it."

Adding to the problem, there is a nationwide shortage of technical security professionals. Private industry has turned to higher salaries and aggressive recruiting efforts to fill gaps. The Federal government's response has been slower and less effective. This means that many Authorizing Officials are not security professionals.

"We have so much complexity and very few people who understand it. The biggest risk factor to me is the breadth and depth of talent we have in the Federal government. It's so hard to attract people, and then even when you can attract somebody, it's just so hard to get them through the hiring process."

"The Marine Corps has one AO right now; one Authorizing Official."

To truly master the Risk Management Framework requires a significant understanding of policy, technology, and the risk landscape. While the technology and security expertise are hard to come by, the Risk Management Framework is readily available. In the absence of other security resources, Authorizing Officials often turn to the paperwork to make decisions.

> *"A lot of them kind of seemed to come up through compliance backgrounds, so they had a decent familiarity with the security landscape, but they definitely didn't come from what I'd consider a real software engineering background. That could make it really difficult to have conversations about certain types of compensating controls"*

> *"Having a script is great. It helps you get repeatability but if you don't understand why things are the way they are, maybe you should not be in that role."*

The Framework documentation is more than 3,500 pages of dense material. It is written at a college reading level or higher. It would take an average person approximately 80 hours to read through the policy and guidance just once. We spoke to several Authorizing Officials, both present and former, that had not read completely through the documents at all. Most had some familiarity with SP 800-53, the list of security controls, but almost no one had read SP 800-37, the Risk Management Framework itself.

> *"This was someone who was extremely well versed in the paperwork of the controls and not really in the underlying technical realities. And so that had a whole bunch of really bad consequences when it came time to decide how to apply flexibility from RMF."*

NIST states clearly that the Framework should not be used as a checklist, but it does so only two times throughout the entire library; 64 words out of 3,500. SP 800-53, the most commonly read and referenced document (see Appendix A), on the other hand, contains a list of more than 1,000 possible security tasks to be accomplished. NIST can say they recommend Authorizing Officials should not use the Framework as a checklist, but, in the words of one official, "then they created all the checklists."

Fear of liability leads to poor decisions

> *"It's a fear based system and no one ever gets in trouble for following the status quo."*

Authorizing Officials believe they are held personally responsible for potential system failures and therefore enforce higher levels of categorization and apply more security controls than make sense from a security perspective. This feeling was more influential on Authorizing Officials than OMB or agency guidance instructing them to use as few controls as necessary to maintain security.

The concept of personal liability was brought up by each person we spoke to, with many saying they or an Authorizing Official they knew required additional paperwork or processes in order to avoid being pulled in front of Congress, being fired, or sent to prison if a system was found to be insecure. This perception was universal, mirrored by policymakers, agency teams, developers, and Authorizing Officials.

A sense of fear creates a desire for self preservation that is not conducive to sound security practices. While people should feel invested in the success and security of the system they are

putting in place, being fearful of losing your job or going to prison puts an unnecessary mental strain on all parties to the process.

Each person we spoke to acknowledged that Authorizing Officials did not want to approve something, have it fail, and be punished. The most commonly cited reason for this fear was the data breach at the Office of Personnel Management (OPM) in 2015, in which stolen employee records compromised the identities of more than 22 million people. In that instance, the OPM Chief Information Officer was asked to appear before Congress, but retired days before she was scheduled to testify. She was never subpoenaed and, ultimately, did not testify. Following several hearings, a Congressional report was released.

> *"It's just the OPM example I'm sure, which is what everyone brings up. And that is a failure of such catastrophic proportions that it is so easy for members of Congress and their staff to understand it. It is truly a unicorn in all of this stuff, right? It is not normal."*

During our research, we were unable to find a single instance of an Authorizing Official being fined or going to prison even in the most dire security incidents.

> *"No one's ever been to FISMA jail."*

That does not mean that fears of routine oversight lack merit, only that the most commonly cited, and extreme, cases are untrue. Congress, either directly or via the General Accountability Office (GAO), does conduct security oversight and routinely relies on the Risk Management Framework to determine if decisions were made correctly.

In addition to Congress and the GAO, interview participants cited inspectors general, The Office of Management and Budget, and the press as forms of oversight that incentivised Authorizing Officials to apply most or all of the security controls listed in SP 800-53 to systems, even when those controls were not applicable.

> *"They're always scared of what their IGs might say or what Congress might say, or what folks at OMB might say, if they take what they presume to be a measured and thoughtful risk or to move quickly in authorizing a new IT system or a cloud service or digital service within their environment."*

Surprisingly, several practitioners expressed sympathy for Authorizing Officials. They noted that AOs were overworked, underprepared, and acknowledged that they were trapped inside the same Framework as system developers and program managers.

> *"The element of fear is there because there's too much to do, not enough money, not enough tech capabilities, no one documents what they do, and the processes are outdated."*

Ultimately, interview participants felt that fear of oversight was both unnecessary and detrimental to security and agencies' missions.

> *"The fear about getting pulled in front of Congress or someone losing their retirement because they would get in trouble for what we had built because of a security incident would happen. That, unfortunately, was kind of the driving incentive for a period of time."*

> *"If executives who are either political leaders or senior career professionals who have worked for 15 or 20 or 25 or 30 years who have moved themselves up to the top ranks of a very large public sector organization, if THEY are scared of a bunch of 25 year olds and a bunch of Congressmen... that's kind of ridiculous in my opinion."*

> *"It's bewildering to me that some thoughtful CIO or CISO or tech executive somewhere in an agency would worry that Congress, which can't even focus on things like the debt limit or appropriations bills or the NDAA for any significant period of time, that they could possibly make some big story or some big example of an IT failure."*

Participants who had served in both policy and oversight roles cited the need for better storytelling from security officials. By leaving the narrative solely in the hands of oversight bodies, they argue, Federal employees face a daunting, and growing, record of negative press that paints them as incompetent.

> *"I think a lot of IT leaders and procurement leaders and security leaders inside agencies are just not good storytellers .... Most of them are not capable of connecting what they do from a technology standpoint, a risk management standpoint, security, or privacy standpoint to why that matters to Congress, to the public, to agency leaders, to folks in the White House or anything else."*

> *"So the public data all lean towards information that signifies that someone did not follow the rules or did not do what they were supposed to do, or did not do what some policy that was written by some analyst at OMB told them to do. And then they do not do a good job of countering that with their own public narrative or their own public data to say, here is what we did and why we did it. So I think they're sort of just operating with a couple of arms tied behind their back when it comes to the prevailing narrative that has been around for decades."*

The unintended consequence of such overly negative oversight is a culture of fear within the Federal security community. That fear changes the risk analysis for individual Authorizing Officials. Fear makes them more conservative and shifts the focus from appropriate security toward self preservation, to the detriment of mission outcomes.

Success is a lack of failure, it's not about meeting mission objectives

The Risk Management Framework puts in place procedures to eliminate as much risk as possible. From the perspective of the Framework, the only success that can be had is if a system does not fail or is not found to be insecure. Success is measured as "not failing."

There are seemingly infinite ways for the Risk Management Framework to fail and only one way for it to succeed. Furthermore, that success is always in jeopardy. It can be proven in the past and the present, but never the future. Because of this approach, and because nothing can move forward without an ATO, mission goals are effectively seen as a lower priority than data integrity.

> *"At one time, I was explaining that the government would want something to fail closed and they were like, 'What? Don't they want their service to operate'?"*

In the area of technical security, perfection is unattainable. Experienced security professionals expect and plan for scenarios in which they lose data or experience performance failures. Success is measured less in terms of completely avoiding incidents and more by how quickly and effectively incidents are remediated.

Experienced security professionals do not seek to prevent all risks. Even if all risks were known and could be mitigated, it is rarely practical. Instead, experienced professionals evaluate risk and manage it proportional to the cost of failure. In theory, this is the value of the Risk Management Framework. In practice, this balance is rarely achieved.

Authorizing Officials use the Framework as a checklist

> *"I've never met [an Authorizing Official] that they're like, 'Yes, my job is to help with security.' It's like, 'No, my job is to make sure the checkboxes are checked.'"*

When asked about the Risk Management Framework, developers, security experts, and Authorizing Officials universally talked about SP 800-53, the list of security controls, or "the checklist." Most had never read 800-37, the Risk Management Framework.

> *"I haven't read it, but as I understand it, the document that establishes RMF is very simple."*

> *"What most agencies that I interact with have done is that they've gone out and, and bought some checklist that they can check off the boxes on and comply with the Risk Management Framework. So, you know, a lot of them may not have ever read the document."*

The National Institute of Standards and Technology has said that the Framework only offers recommendations, not mandates. As previously discussed, these recommendations are routinely

interpreted by the Authorizing Officials and security professionals as checklists that must be completed for an authorization to be granted.

> *"I hate to say this so insultingly, but you can't just wave a magic wand and have the security staff actually know things about security when their jobs have been doing paperwork checklists for decades."*

Using checklists is a reasonable approach given the circumstances. If Authorizing Officials are not technical enough to engage in meaningful security discussions with developers or if they do not have the time to dig into system details, the Framework provides them with a proxy method for completing security processes. Given both the real and perceived consequences for failing to follow the security controls, it is reasonable for Authorizing Officials to adhere strictly to the full list of security controls.

> *"RMF is very checkbox focused and not security focused"*

> *"We shouldn't be spending money to try and mitigate risk that doesn't exist."*

Approval is relational

The security of Federal technology often comes down to factors that are unrelated to technical risk or system integrity. Poor guidance, a lack of knowledge and training, time constraints, and unintended incentives lead Authorizing Officials to make subjective decisions about risk. Many Authorizing Officials use trust or other factors as a proxy for security expertise.

> *"It comes down to your individual security organization and your authorizing official and what they feel comfortable signing off on."*

> *"That's why the process is highly relational. The guy who ended up signing it is not somebody who was very good at cybersecurity. He just said 'Okay, we've built up enough trust that, finally, I will cave and sign this for you.'"*

> *"It's very personality driven, especially on some of the more sensitive controls."*

> *"The idea around the authorizing official is having someone who has some skin in the game, who is going to be providing this oversight, and who's going to be hopefully above any conflict of interest to say, 'Wow, that's not secure, let's turn that off.' Or, 'We need to apply these resources to fix this problem.' And that doesn't really happen."*

> *"Typically it's saying the right word, finding the right people, and hoping that you get through in a reasonable amount of time."*

> *"It depends on the people you're dealing with in terms of how open minded they are and what you say to them."*

While most interview participants highlighted improving trust and interpersonal skills as the most effective path to approval, Authorizing officials used any number of requests to ensure they were comfortable approving and authority to operate.

> *"The alternative implementation from [the Authorizing Official] basically amounted to encasing the cable in a foot and a half of concrete."*

> *"One of the security engineers we're talking with suggested gluing shut the ethernet jacks on the laptops with epoxy to prevent them from wiring into the network because there were no cryptographic controls on the [agency] network at that time."*

> *"And this goes back to the ATO being highly relational. We got it signed on my birthday because we had a meeting with the signing official the day before my birthday. And I said tomorrow is my birthday and I would like this to be signed."*

While some requests seemed unreasonable at the time, interviewees recognized that, at the end of the day, the Authorizing Official could prevent systems from going live. If an Authorizing Official were adamant about a process or mitigation, the development team would have to meet their requests.

> *"They're not bad people. They just get bad pressure. It's not a psychologically safe org, the government. It's not a very good place for them to say, 'I don't know,' or 'maybe we could do that.' All of those things introduce danger for them. To put their neck out, there's low, low incentive for them to do that."*

> *"Ultimately it's your AOs butt on the line, right? They're the one who's going to have to go testify to Congress if your app gets hacked. So whatever they need to see in that ATO package to make them feel comfortable signing off on the risk is what needs to be in there."*

> *"From their perspective they're acting totally rationally. Their incentives are purely to try to stop things and delay them as long as possible. There's no benefit to them in allowing a new thing to go forward."*

Changing incentives for Authorizing Officials change the risk calculation

Federal agencies have to contend with a landscape that is constantly changing. When people come and go from security roles in an organization, the risk tolerance can change. Knowledge may be lost or there may not be sufficient documentation to understand why prior risk decisions were made. Time may simply have passed or the underlying reasoning may no longer apply.

This is especially true when Authorizing Officials change. The risk calculation of a previous Authorizing Official may not be in line with that of the new official. Not only is there irregularity and inconsistency in the way systems are authorized, there is irregularity in the way they are monitored, managed, and reauthorized over their lifespan.

> *"We started the project with maybe two to three months of high level air cover... and then as soon as the transition happened, [the Authorizing Official] was out. We had a new CIO who could not actively kill the project but was not just not at all interested in lending their support."*

> *"We have our entire public cloud in fedRAMP Moderate .... We do have times where we have to go back to the Authorization Board or go back to the Authorizing Official and talk through why the previous staff agreed to this and fill them in in terms of what we're doing. And every now and then a different interpretation results in changes that we have to implement to maintain compliance."*

This shift in incentives can also occur after systems are authorized for the first time. Before a system is authorized, the incentive structures encourage the maximum number of controls. They effectively discourage authorization. Waivers are difficult to get. Once a system is authorized, and after it becomes part of the way the agency works, that dynamic shifts. It becomes difficult not to re-authorize a system that is already in use. Waivers become more frequent. Still, the incentives are not necessarily to modernize or improve the system. Too many changes would trigger a new ATO process. Therefore, when reviewing the security of an existing system, the incentive is to change as little as possible so the existing authority to operate remains valid.

> *"Same thing with a CIO where if they have an old system from a previous CIO that the previous CIO approved, they have the option of improving it and changing it up, and therefore their signature has to go on all the new documentation, or leaving it as it is. Then if it breaks, it's not their fault, it's the fault of the previous person that approved it."*

Developers and Program Officers must contend not only with the personalities and risk tolerance of one Authorizing Official. Over time, they are likely to encounter several officials with different opinions and comfort with risk. This is especially problematic for private sector vendors who sell off-the-shelf solutions to the government. The more variability they encounter, the less attractive the government is as a market.

Vendors

The Risk Management Framework governs more than just proprietary government software, it applies to all commercial software used by the government. Software vendors play a critical role in government. Software vendors can either work directly with agency Authorizing Officials to integrate their products into government work environments or they can take advantage of programs like FedRAMP, which are designed to pre-approve security compliance for some or all

aspects of their products. We spoke to commercial vendors attempting both approaches and found several common themes.

Decision making is inconsistent

Each Authorizing Official can, and often does, set different security requirements. Some Authorizing Officials may expect exact adherence to every control while others may require additional overlays or bespoke changes. Vendors struggle with this lack of consistency. One large technology company we interviewed said that they were managing multiple Authorities to Operate, implementing different requirements, at the same agency, for the same product, even though that product had already received FedRAMP approval.

> *"FedRAMP was supposed to just be like, 'We're good to go.' But no, that's not at all how it happens in practice."*

> *"We had to suddenly implement things that the commercial customers were like, 'Oh, no.' The 15 minute logout.... I would've been ridiculed by my own engineering team if we had pushed that through for our own staff."*

> *"I totally empathize with the agencies, like they're just trying to meet their requirements and make sure that they're shored up... our customer wants to know if they can check the box and say they meet a specific control with our software and I'm like, 'Okay. Well, it's not a quick ask.'"*

> *"I mean hopefully our little written memos stick and we can just continue doing business with some level of assurance that we won't be getting requirements changed on us and investments will be lost in the future."*

The Framework does not align with how technical vendors work

Commercial products that are submitted for ATO approval cannot be updated easily. When a new feature is commercially available, government customers receive it only after it has been reviewed. In some cases, there is additional delay while code is transitioned into government specific infrastructure. There are even times when the government does not immediately benefit from security updates. Furthermore, there are some features that will never be made available to the government because of the burden of security compliance.

> *"Right now I hear that it's six months before you can get an audience with the Program Management Office."*

> *"If you're really concerned about a wartime scenario, then it just makes sense that you would want to be able to have redundancy. So some of those threat models when we walk through with them, they're not necessarily able to articulate the problem that*

they're trying to solve and we have to say, 'Well then we don't know how to solve it for you.'"

The cost of working with the government is high

To navigate the Risk Management Framework, vendors must answer an Authorizing Official's questions about their commercial security management practices. Vendors spend significant amounts of time, effort, and money completing the ATO process. At times, vendors said they needed to reduce the security of their product or eliminate features altogether in order to receive authority to operate.

> *"As soon as you go into the government space, it seems like you pay 30% more."*

> *"Oh definitely! We have to disable features, for sure, based off of the compliance requirements."*

To work with agencies that require the highest FISMA or FedRAMP approvals, vendors often have to make difficult decisions. Some decide to divide their companies (or product lines) into two organizations with one focused solely on meeting government-specific requirements. Others decide simply not to sell to the government. This is especially true of small businesses who cannot shoulder the cost or commit the dedicated time to meet the Framework requirements.

> *"I think we would need to split into a commercial deployment and a Federal deployment and build a moat around it and fill it with laser sharks. And I'm not sure if that's something we have the will to do."*

To be successful in the healthcare space, one company we spoke with felt they needed to work with the Department of Veterans Affairs (VA), the largest health provider network in the United States. Due to the VA's FedRAMP High requirements, however, the company could not justify the enormous burden and cost of getting and maintaining their security compliance requirements. This left the company and their clients, including the Department of Health and Human Services, with significant gaps in their capabilities.

The Framework reduces competition & offers the government fewer options

In practice, the RMF process creates an environment in which the government simply does not have access to many technologies. The cost and complexity associated with the Framework and programs like FedRAMP unintentionally reduce competition. When vendors are willing to work with the government, the Framework increases prices and leaves the government with fewer tools and features at its disposal.

"The real question is, is the Federal government better off for restricting its suppliers to the ones that can afford to throw down a couple of million dollars beforehand?"

"[Leadership] told us we were never getting this FedRAMPed because FedRAMP cost us $500,000 to two million dollars. We're not going to do that to sell a ten dollar subscription per month."

Recommendations

Write "This is not a checklist" on every page of the framework

The most common misunderstanding in the risk management process is that the Framework is seen by many as a checklist of security tasks rather than a menu of potential security issues to consider when authorizing a system. NIST and OMB both state that agencies should tailor security requirements to those that are relevant to a particular system and thread environment. They both state that mission needs and cost should be considered when managing risk. Both sentiments, however, are lost in the lengthy, dense, and bureaucratic Framework literature.

If NIST believes that this misunderstanding is problematic, they could reinforce their guidance by repeating it in every document they produce. SP 800-53 implores readers not to use the controls as a checklist, but that message has not broken through and is therefore worth reiterating on every page. Perhaps because most practitioners do not read SP 800-53 as guidance, they treat it as a reference manual, skipping to the relevant controls to better understand specific details rather than treating it as a holistic narrative. NIST should reinforce the behavior they want Authorizing Officials to take on every page of SP 800-53 and, if necessary, the entire Framework.

Shorten the Framework & use plain language

The Risk Management Framework should be shortened dramatically and re-written in plain language. As it is, the Risk Management Framework is difficult to understand and navigate. Its sheer size is intimidating. Iit requires at least a college reading level to understand, to say nothing about the technical knowledge required to correctly implement it. Despite this, the Framework documents are meant for a broad audience, including many non-technical senior executives, program officers, and civil society and private sector partners. Plain language would also make it possible for members of the general public, including the media, to understand how the government manages technical risk and hold it accountable.

Unfortunately, despite Federal plain language guidance and best practices, the Framework remains out of touch for many key stakeholders. We interviewed more than 20 practicing security professionals who, despite their reliance on it, had read few, if any, of the documents. Even fewer took advantage of the flexibility and adaptability allowed by the Framework. Those who did rarely encountered counterparts, developers or Authorizing Officials, with a similar level of understanding.

In the rare case of projects in which program managers, developers, and Authorizing Officials all collectively agreed on the intent and flexibility of the Framework, authorities to operate were approved in as little as hours or weeks.

The relative inaccessibility of the Risk Management framework is a security risk. When experienced professionals cannot, will not, or do not have time to read policies and guidance, they cannot implement it successfully. If NIST wants teams to benefit from the Framework, they need to understand who their readers are and how they are using the documents. If they understand their audience, they can write in a way that is inclusive, accessible, and useful. People will not follow guidance they cannot understand.

Update A-130: Make it actionable, readable & prominent

The OMB guidance for the "Management of Information as a Strategic Resource" is six years old, far too outdated in an era of rapid technical change. OMB should update A-130 to align it with current priorities and technologies, such as the recent cybersecurity and customer service executive orders, the findings of the equitable data working group, the AI Bill of Rights, as well as modern cloud and artificial intelligence capabilities. In addition, OMB should work with NIST and the CIO counsel to align agencies around common talking points regarding technical risk management.

While there is good guidance in A-130, it is neither prominent nor actionable. It is buried in definitions, appendices, and footnotes, making it difficult to find and use. OMB should highlight best practices in security more directly and in plan, directive language. If security management has strayed from the intent of NIST and OMB, OFCIO should address it directly and outline acceptable alternatives. OMB could also use A-130 to create a permissive structure for Authorizing Officials to make decisions that are in the best interests of the mission of government, shielding them from undue reprisal if they act in the national interest.

A-130 could also be used to set professional standards and create requirements or certification credentials for Authorizing Officials. This would ensure a baseline of knowledge for Authorizing Officials, greater consistency in the application of security processes, and improve discourse between security officials and private sector vendors.

Lastly, A-130 is difficult to find, even on the official cio.gov website. The link from the Federal register is no longer active and the secondary link on the OMB website links to an archived version of A-130, which appears to be the most current. The age of the document and the lack of a current copy on official websites, diminishes the authority of the guidance.

Update documents more frequently

Software security is a rapidly evolving field that requires continuous updates and evaluation. Policies should be kept current and evaluated regularly.

The current iteration of the RMF began in 2004 with the deprecation of FIPS 102 and the introduction of FIPS 199 and SP 800-37. FIPS 199 has not been revised since 2004. 800-37 was not revised until 2018.

Older documents can create conflict with modern policies. For example, portions of FIPS 200's Minimum Security Requirements (written in 2006) could be construed to prohibit or limit efforts to move toward zero trust environments, an administration priority. This conflict was unavoidable at the time FIPS 200 was written, as zero trust concepts did not emerge until 2010. FIPS 200 has not been updated for seventeen years, creating confusion as to how the two policy objectives might interact.

Documents that are three years out of date raise suspicion in the minds of many security professionals. RMF documents, on average, are nearly nine years out of date. We recommend that NIST review documents every two years, three at maximum. When they do not require substantive changes, NIST could mark them prominently as current to reflect their continued relevance. This could include putting the new document version checker button[5] on each document to let the public know it is current. More frequent updates will increase confidence in the relevance of RMF documents.

Professionalize the Authorizing Official role

Professionalizing Authorizing Officers is the most effective way to improve both the efficiency and effectiveness of Federal technical security within the current Framework.

Authorizing Officials play an equally important and critical role as more officially accredited officials in the Federal government, such as Procurement or Consular Officers. They are responsible for making a final yes or no decision about issues of national security. They bear accountability and incur some liability, if not as much as they perceive they do. Despite that, there are no clear roles and responsibilities that are assigned to all Authorizing Officials. Training is not formalized and there are no professional certifications or warrants.

Without a clear position description, there is variability in the skills and experience of Authorizing Officials. Some are experienced security professionals while others are tasked with the authorizing responsibility as a secondary duty, sometimes outside of their primary profession. Inconsistency undermines the credibility of qualified officers and instills overconfidence in unqualified officers. This lack of support creates variability in authorizing decisions. It also fosters distrust among Authorizing Officials which undermines the proper scoping of boundaries, tailoring of controls, and reuse of common controls.

[5] New NIST documents have a tool embedded into them to help readers ensure that the documents are current. Currently, only two Framework documents have this helpful feature.

Setting a minimum standard for Authorizing Officials, in addition to subject matter-specific responsibilities, would ensure more reliable, replicable outcomes. Instead of enforcing the entire, or nearly entire, set of security controls as a means of limiting their liability, officers could point to training materials and instructions for why they made decisions. Proper training, certification, and ongoing professional development would also allow the government the ability to update training materials to meet current advances in security best practices. Regular training would give the government the ability to adapt implementation practices much faster than updating the Risk Management Framework library.

Consider narrative security approaches

The Federal government could adopt a narrative approach to security that is more accessible and understandable than the traditional ATO process. Doing so would effectively encourage mission owners, developers, and security professionals to discuss both system security and understand reasonable steps to be taken before, during, and after any potential security incident. Good examples include the United Kingdom's basic risk assessment and management method and lightweight approach to cloud security. Both are effective at reducing the complexity of security assessments while encouraging a dialogue that ensures teams default to modern security practices.

This type of approach builds on the intent of agency and FedRAMP "lightweight" processes. Narrative ATOs build off of the idea of attesting to routine security processes, and focusing more rigorously on key areas. Further simplification and plain language would improve communication and include more stakeholders in security discussions.

Narrative ATOs can be used in conjunction with or as a proxy for the control list in SP 800-53. For less technical practitioners, it can help teams focus on areas of concern, at which point security professionals could apply existing controls. Ideally, however, this approach would replace large sections of the Framework library and simplify ATOs to a point that they could be completed, read, and understood by both seasoned security professionals and non-technical project managers. A lighter-weight approach to the security process would also be faster, helping the government keep pace with technological innovation.

Do away with all or part of the Risk Management Framework

20 years ago, the risk management process made technical security a priority at a time when it was not commonplace to consider the risks associated with technology. Today, however, technical security is a well-known, if sometimes misunderstood process. Without a formal framework, security experts and developers today would still consider system security as part of their professional responsibilities. Agencies would still develop authorization procedures tailored to their unique risk posture. Oversight bodies would still hold the government accountable for failures to protect Federal data. NIST would still play a vital role in the strategic understanding and categorization of risk. They would all simply do so without the burden of an extensive and outdated

compliance process. In short, they would not be beholden to the checklist mentality or the categorization trap if there were no checklist or categories.

If the goal of the Framework is to raise awareness of security issues and begin an honest dialogue between security experts and development teams, a much simpler, less prescriptive process would be more effective and much more likely to become a routine part of the development of products, processes, and policies.

Conclusion

When Congress wrote, and later updated, the Federal Information Security Modernization Act, it could not have known all of the ways it would affect security in the government. Clearly, its intentions were to improve data security across the enterprise and make it easier for the private sector to work with the government. The resulting Risk Management Framework and policy guidance were also well-intentioned and comprehensive. NIST sought to understand all possible risks and provide guidance that everyone could use.

NIST intentionally created the broadest understanding of risk they could, leading practitioners such as the MITRE Corporation to the conclusion that: “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.” To counter that breadth, NIST built a great deal of flexibility into the Framework and wrote a series of guides dedicated to helping practitioners apply key areas of the policy. OMB followed-up with practical guidance for Federal agencies that encouraged a balanced approach to risk that took into account cost and mission outcomes.

For their part, Federal agencies developed authorization plans and designated Authorizing Officials to implement the Framework in the context of the agency’s mission and risk posture. Agencies designated Authorizing Officials who took on the task of managing risk as best they could. They created individualized risk assessments and plans for almost every technical system in government.

Vendors and development teams worked together with Authorizing Officials and program offices to try and understand and manage risk. Everyone in the process took rational steps within the confines of the Framework and the incentive structures put in place by the government. It did not lead to efficient security outcomes.

The policy does not translate into effective risk management. Instead, the Framework leads to risk avoidance, both personal and professional. Without any entity acting inappropriately, the complicated incentives of government led to a process that was slower, costlier, and less effective than intended.

The Framework increased reliance on legacy systems and reduced the number of commercial solutions available to the government. In some cases it degraded security and contributed to an increased likelihood of potentially harmful events, such as the inability to deliver key public and economic services. This was an outcome that no one wanted or foresaw, but it has become widely accepted that the Framework is too slow and cumbersome to be an effective tool for managing technical security. The Framework is also too established to abolish. Therefore, technical security has become a step separate from security compliance.

Precisely because everyone in the process is acting honestly and in their own best interests, changing the system is incredibly difficult. Many actors with differing, often competing, incentives

must choose to act in concert and against their best interests, in order for real change to happen. Most of our interview participants found this unlikely or impossible. Few could articulate a way to improve the Framework and many felt that the government would encounter serious failures before the system could change.

Furthermore, in the absence of change, the inherent problems with the current system continue to worsen. Modern technology is increasing in complexity while it is also becoming easier to use. Technologies such as large language models are difficult for even seasoned technical experts to fully understand but make it simple for non-technical users to complete complex tasks, such as writing code or querying large data sets. As more people use more complex technology more frequently and in more contexts, NIST must consider new risks and incorporate them into the Framework. Authorizing Officials will need to manage more technologies and authorize new systems while maintaining an increasingly large portfolio of ongoing monitoring and reauthorization.

In its current state, the Framework will continue to grow in scale and complexity. New controls and overlays will be necessary to manage new use cases. The introduction of the new AI Risk Management Framework is a good example of how the paperwork struggles to keep pace with technology. This will increase the time and costs associated with authorizing new technology, exacerbating the fundamental flaws of the Framework.

There are short term approaches that would improve the implementation of the current Framework. A professional class of Authorizing Officials, with relevant expertise, adequate training, and reduced liability would deliver better results. A simplified ATO process focused on outcomes would help encourage security over compliance. NIST could narrow its focus to the most critical systems or write for a targeted audience, such as inexperienced practitioners or owners of legacy software. Across all parties, clearer communication and plain language would encourage greater understanding, clarify intent, and improve the quality of security dialogue.

Ultimately, however, the system itself must change. Managing security via paperwork and personalities can neither capture the dynamics of technical security nor can it keep pace with technological innovation. If technical security is to remain a human process, it must be drastically simplified so that its users can keep pace with current technical needs; pruned down to its essentials and re-written in clear, direct language. Culturally, the government must learn to accept greater risk and elevate the risk of mission failure as a key factor in technical security decision making. Failing to deliver benefits and services, in most cases, will lead to greater harm than the loss of data, especially when appropriate safeguards and failsafes are put in place.

The Framework is well-intentioned but it has become increasingly cumbersome and ineffective at managing security as technology has evolved.

Acknowledgements

Thank you to the past and present staff members from the Executive Office of the President, the Senate, the Department of Homeland Security, the General Services Administration, the Department of Veterans Affairs, the Department of Health and Human Services, The Department of Commerce, The National Aeronautics and Space Administration, the National Institute for Standards and Technology, and the many others who supported this research with their time and input.

Thank you to our private sector experts for your perspective and insights into working with the Federal government as well as your independent perspectives on technical security.

Thank you to our early readers for their help in making this report comprehensive, readable, and accurate.

Thank you to our philanthropic sponsors for their generous funding and support of this work.

Appendix A: RMF relationships between documents

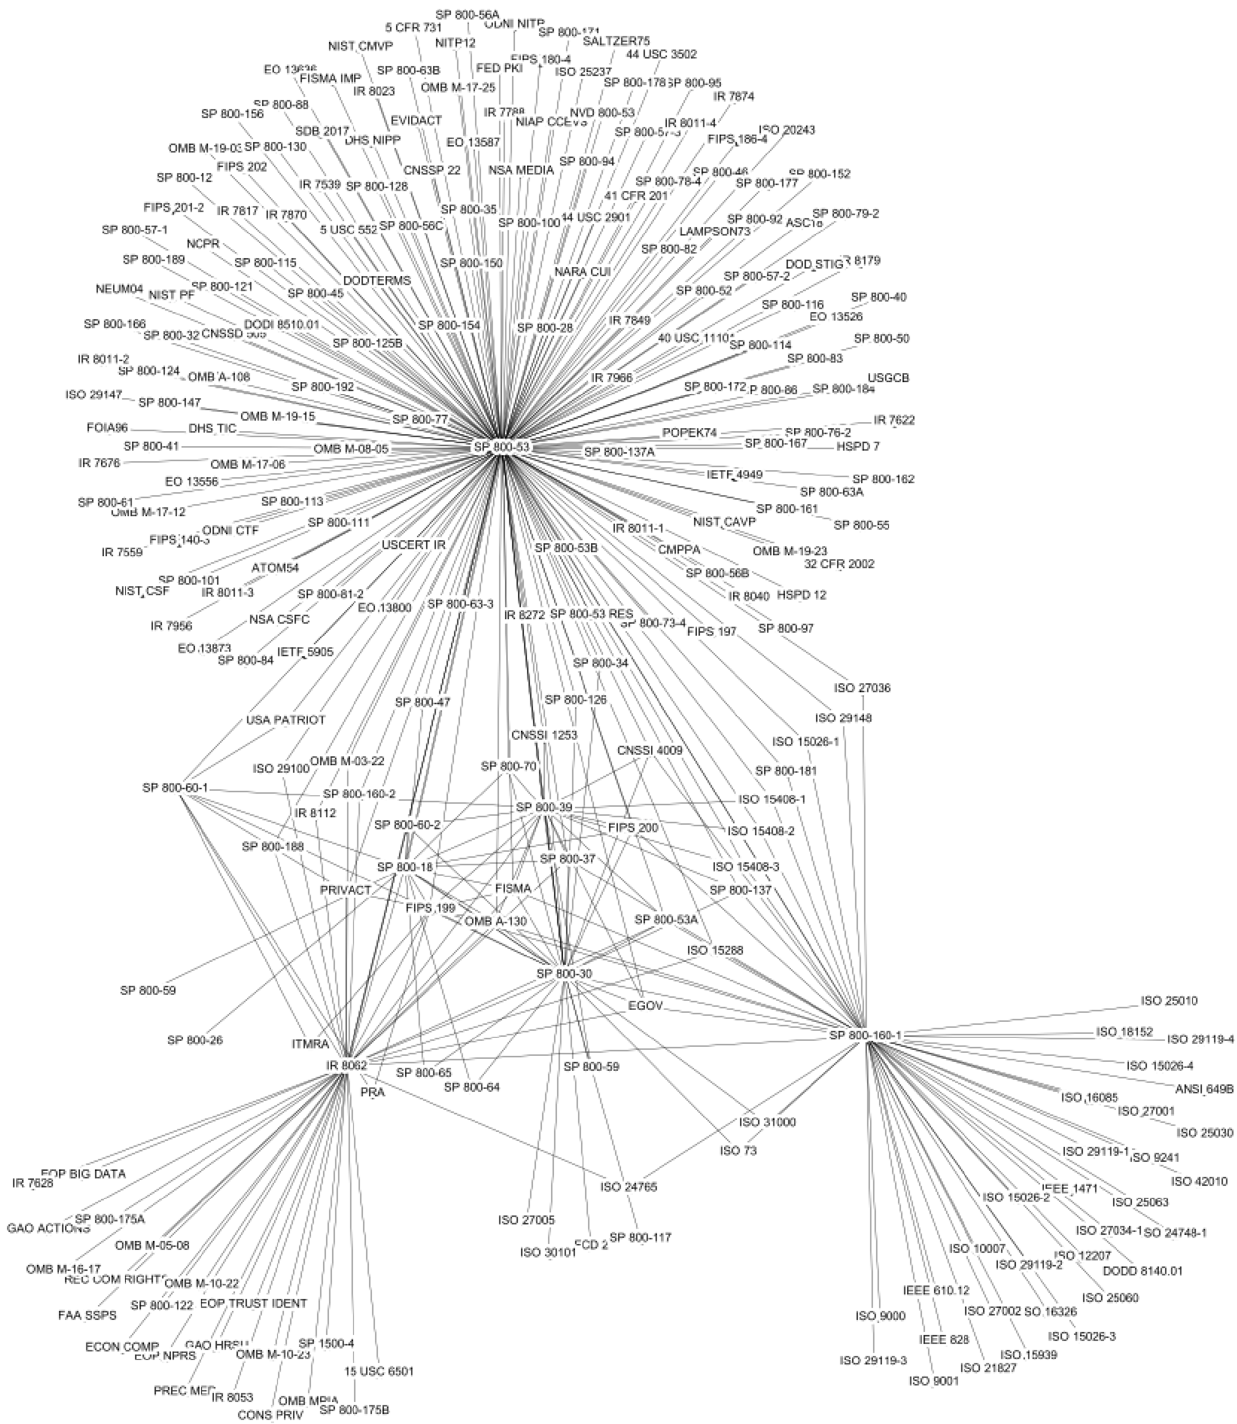

Appendix B: Documents reviewed

Stage	Title	Document number
Overview	Risk Management Framework for Information Systems and Organizations	SP 800-37
Prepare	NIST PRIVACY FRAMEWORK CORE	NIST Privacy Framework
	An Introduction to Privacy Engineering & Risk Management in Federal Systems	NISTIR 8062
	Engineering Trustworthy Secure Systems	SP 800-160v1
	Developing Cyber-Resilient Systems	SP 800-160v2
	Guide for Conducting Risk Assessments	SP 800-30
	Managing Information Security Risk	SP 800-39
	Guide for Developing Security Plans for Federal Information Systems	SP 800-18
Categorize	Standards for Security Categorization of Federal Information & Information Systems	FIPS 199
	Guide for Mapping Types of Information & Information Systems to Security Categories	SP 800-60v1
	Appendices to Guide for Mapping Types of Information & Information Systems to Security Categories	SP 800-60v2
Select	Guide for Mapping Types of Information & Information Systems to Security Categories	FIPS 200
	Security & Privacy Controls for Information Systems and Organizations	SP 800-53
	Control Baselines for Information Systems & Organizations	SP 800-53B
Implement	Contingency Planning Guide for Federal Information Systems	SP 800-34
	Computer Security Incident Handling Guide	SP 800-61
	Guide for Security-Focused Configuration Management of Information Systems	SP 800-128
Assess	Assessing Security & Privacy Controls in Information Systems & Organizations	SP 800-53A
	Automation Support for Security Control Assessments	NISTIR 8011
Authorize	Developed by agencies	varied
Monitor	Information Security Continuous Monitoring (ISCM) for Federal Information Systems & Organizations	SP 800-137
	Assessing Information Security Continuous Monitoring (ISCM) Programs: Developing an ISCM Program Assessment	SP 800-137A
	ISCMA: An Information Security Continuous Monitoring Program Assessment	NISTIR 8212
Overlays	Security Control Overlay of NIST Special Publication 800-53 Revision 5 Security Controls for Federal PKI Systems	FPKIPA
Other	Managing Information as a Strategic Resource	A-130
	Framework for Improving Critical Infrastructure Cybersecurity	Cybersecurity Framework

Appendix C: Reports & references

- NIST working on 'potential significant updates' to cybersecurity framework
- Agencies Need to Develop Modernization Plans for Critical Legacy Systems
- Bill would reform cybersecurity management
- Authorization to Operate Field Guide
- Oversight.gov
- Plainlanguage.gov
- How Complex Systems Fail
- Conway's Law
- Research: Simple Writing Pays Off (Literally)
- OPM Cybersecurity Incident
- General Accountability Office: Healthcare.gov
- Cybersecurity & Infrastructure Security Agency: Cybersecurity/IT Careers
- The Cybersecurity Workforce Gap
- Beyond Compliance–Addressing the Political, Cultural and Technical Dimensions of Applying the Risk Management Framework
- Federal IT Dashboard
- Federal Agencies' Reliance on Outdated and Unsupported Information Technology: A Ticking Time Bomb
- Agencies Need to Develop Modernization Plans for Critical Legacy Systems (2019, 2021 & 2023)
- Shadow IT Provides Clues to the Tech That Federal Workers Really Need
- GSA security authorization guidelines
- GSA Lightweight Authority to Operate process
- FEDRAMP tailored authorization for Software as a Service (Li-SaaS)