Policy Brief

Outpaced

AI and Policy's Role in Transforming Cybersecurity Compliance

Author
Katherine Carroll

CSET CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

August 2026

Executive Summary

American commercial industry leads the world in software development and artificial intelligence (AI). That advantage does not automatically transfer to the military. The federal cybersecurity compliance process is a major barrier between America's most advanced technologies and the warfighters who need them. The Authorization to Operate (ATO) is the federal government's formal mechanism for assessing and approving software systems. However, the ATO process and its governing framework—the Risk Management Framework (RMF)—are described as ineffective, slow, duplicative, and in need of reform.[1] Bureaucratic delays can carry a devastating cost. A former intelligence officer describes just how high those stakes can be: "I firmly believed that software was the reason that a bunch of *civilians had died*. . . . We had critical fixes like known operational issues that were causing operational risk that were sitting on the shelf waiting to go through the ATO process."[2]

Prior research has documented ATO process inefficiencies within specific service contexts and cataloged security shortfalls. This paper is the first analysis to examine not only the process itself but the foundational legal authorities, competing stakeholder incentives, and governance structures that collectively produce delays. This paper also assesses why reforms have not solved these issues, despite sustained attention over more than a decade. A summary of the analysis is captured in Figure 1.

Figure 1. Summary of Analysis

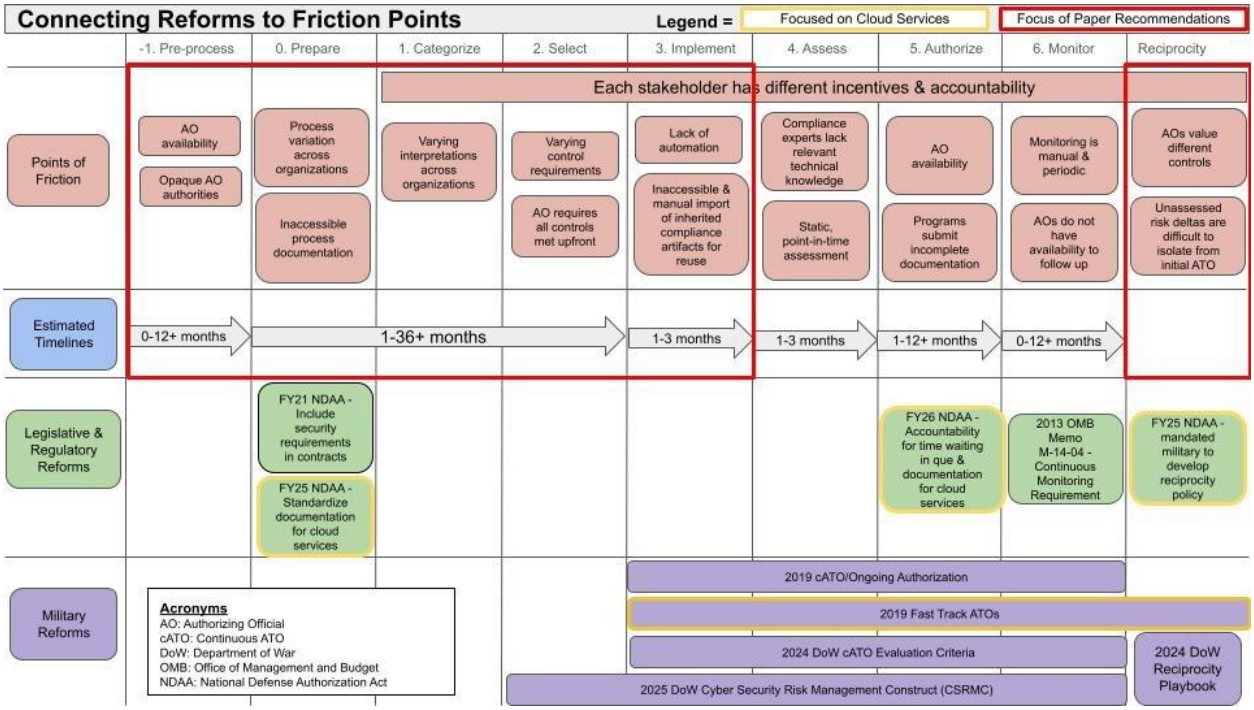

Source: Author's analysis based on 2013 OMB Memo; 2024 GAO *Critical Cybersecurity Challenges;* 2019 *GAO Cloud Computing Security;* 2024 *GAO Cloud Security;* 2024 *DoW Cybersecurity Reciprocity Playbook;* 2021 *GAO Weapon Systems Cybersecurity;* 2016 NIST *About RMF;* 2021, 2023, 2025 and 2026 NDAAs; 2024 cATO Evaluation Criteria; 2019 AF *Fast-Track ATO;* 2025 *FedRAMP Built a Modern Foundation;* 2025 *DoW Cybersecurity Risk Management Construct* and practitioner experiences.

The first half of the military software approval process remains almost entirely unaddressed, and it is the most significant time suck for new commercial entrants. Reciprocity, which allows reuse of ATOs to help reduce duplication, has fallen short of its intended utility and remains a barrier to scaling across the services. The recommendations in this paper identify the highest-return opportunities available that complement rather than duplicate existing reform efforts. These are not exhaustive recommendations for every single point of friction, but a prioritized set of interventions designed to produce the greatest impact given current AI capabilities and policy reform momentum.

Recommendation 1: Publicly release critical controls and authorizing officials. Federal organizations should publish their AO authorities and critical controls to help programs and vendors identify an appropriate AO with bandwidth to evaluate the system. This does not require revealing individual identities, system vulnerabilities, or specific implementation details. Following Karickhoffs's principle, publishing critical controls establishes clear public standards without exposing vulnerabilities, and

publishing delegated authorities would allow AI to assist with AO identification. Opacity creates administrative drag without enhancing security, and administrative barriers discourage the commercial innovation the military needs.

Recommendation 2: Leverage AI to standardize to machine-readable formats. AOs should be required to feed process standards and documentation templates into a centrally managed, AI-enabled tool to standardize ATO submissions. This tool would normalize RMF forms and control taxonomies, assist in generating contract language, and enable automated form population. The military and intelligence community (IC) should also adopt Open Security Controls Assessment Language, the machine-readable format developed by the National Institute of Standards and Technology (NIST) that renders compliance documents as structured, interoperable data. OSCAL establishes a standardized output format that accommodates variation in underlying tools while ensuring that compliance artifacts are interoperable.

Recommendation 3: Budget and deploy continuous, automated AI red teams. Programs must incorporate the cost of AI red teaming into budget plans through a dedicated percentage tax on overall budget dollars or by requiring program managers to explicitly include estimates within formal program estimates. Once a system meets critical security controls, program offices and AOs should stop treating the remaining NIST controls as a compliance checklist and instead implement continuous, automated AI red teaming through the system life cycle. Agentic AI can execute realistic, on-demand threat assessments continuously and at scale, providing real-time security insights while optimizing the deployment of scarce, highly skilled human cyber talent. Currently, the scarcity of personnel possessing deep technical expertise, mission context, and RMF mastery creates cybersecurity compliance theater and delays.

Recommendation 4: Experiment with AI reciprocity agents. An AI reciprocity agent would continuously monitor and pull new assessment reports to feed AI red teams, shifting from static, point-in-time assessments to more timely threat information. This capability would introduce critical security risks—primarily the data aggregation threat—requiring a joint military, IC, and federal AI task force to govern and audit the agents with strict guardrails. This recommendation should not be implemented until there has been a proof of concept with low-impact, unclassified systems to validate accuracy and security guardrails. Systems such as Nuclear Command, Control, and Communications (NC3) should be permanently excluded.

Recommendation 5: Increase reciprocity and cybersecurity incentives. The U.S. Department of War—also recognized as the U.S. Department of Defense—and IC chief information officers (CIOs) should mandate automatic reciprocity for defined system

classes that meet standardized critical controls, shifting the default posture to a "presumption of adequacy" backed by continuous AI red teaming and an explicit "no list" for specific high-risk platforms or missions. A valid ATO issued by any DoW or IC AO would grant immediate, service-wide fielding rights. The Defense Innovation Board (DIB) previously recommended skipping pre-authorization compliance entirely, but a system that was never assessed and subsequently compromises sensitive data cannot be rectified through financial penalty.[3] Instead, the DoW and IC should set clear expectations through published critical controls, establishing an organizational bias toward rapid scaling of technology while excluding mission systems such as NC3, next-generation fighter aircraft, and nuclear submarines.

Finally, the military and the broader federal government need a more nuanced approach to cybersecurity incentives than the binary ATO status quo. Contracting mechanisms such as award feeds could incentivize contractors to respond rapidly to relevant threats discovered by AI red teaming. AOs and CIOs could establish guardrails that automatically remove systems that fail to prioritize and resolve known vulnerabilities. Contractors and program managers could receive cybersecurity report cards that follow their professional records, creating accountability that persists beyond any single program.

These recommendations carry real risks that must be managed: Aggregation of vulnerability data, automation bias, cognitive off-loading, and AI accuracy limitations all require deliberate governance structures, phased implementation, and continuous human oversight. The current ATO system already accepts significant risk. Assessment personnel rarely possess deep technical expertise in the latest innovative commercial software, the threats to that software, mission context, and RMF mastery simultaneously, which produces compliance theater rather than real security. AI red teaming and reciprocity agents offer an opportunity to more actively manage the already widespread risks. By transitioning the compliance workforce from attempting to master every technology and threat to governing and auditing AI red teams and agents, the government can scale scarce expertise in ways that human staffing models cannot. Some personnel will still need deep technical knowledge to train models, validate outputs, and build guardrails so that the red teams and agents do not turn against the mission systems. But the alternative of continuing to accept the compounding risks of an under-resourced, checklist-driven compliance system is not a winning risk management strategy.

Table 1. Summary of Recommendations

#	Recommendation	Key Actions	Impacts	Policymaker	Difficulty
1	**Publish critical controls and authorizing officials**	Mandate **public release** of AO-delegated authorities and critical controls.	AO matching is faster and increases the likelihood that the AO understands the technology. Schedule and cost risk is decreased for contractors and program managers since the compliance expectations are clear.	OMB, Congress, or DoW and IC CIOs	**Easy**
2	**Leverage AI to standardize to machine-readable formats**	Require **OSCAL**-formatted compliance outputs across all ATO submissions. Deploy **AI-enabled tools** to normalize documentation and auto-generate artifacts.	Reciprocity and automation increase. Competition of DevSecOps and compliance tools remains, but compliance is automated for auditors. Auditors are not required to be experts in every tool or technology and mission context.	OMB, Congress, or DoW and IC CIOs	**Medium**
3	**Budget and deploy continuous, automated AI red teams**	Require budgeting for and deployment of **AI red teams** for continuous threat testing across the entire system life cycle.	AOs and assessors can implement more effective threat monitoring without having deep expertise in every threat or technology. Program managers, contractors, and operational users can better prioritize the security backlog based on threats and mission context.	OMB, Congress, or DoW and IC CIOs	**Hard**

#	Recommendation	Key Actions	Impacts	Policymaker	Difficulty
4	**Experiment with AI reciprocity agents**	**Connect federal authorization repositories** to allow AI to retrieve appropriate assessments for reuse.	AI agents would have more up-to-date threat information, vast expertise on various types of technologies, and mission-specific context.	OMB or Congress	**Hard**
5	**Increase reciprocity and cybersecurity incentives**	Establish presumption of adequacy for systems with an ATO scaling across the DoW. Leverage other financial and performance incentives outside of the ATO process to help prioritize cybersecurity.	Technology scales rapidly across the military after initial ATO is achieved. Programs and contractors prioritize cybersecurity features due to the apparent fiscal impacts.	Congress or DoW and IC CIOs	**Easy**

Table of Contents

Introduction

Barriers to data sharing, authority to operate, or ATOs, test and evaluation, and contracting are now treated as operational risks, not simply bureaucratic inconveniences.

- Secretary Pete Hegseth, 2026

Modern warfighting is increasingly defined by software that fuses intelligence, enables command and control, and allows military forces to adapt faster than their adversaries. But even though American commercial industry leads the world in software development and artificial intelligence (AI), that advantage does not automatically transfer to the military. The primary barrier between America's most advanced technologies and the warfighters who need them is the federal cybersecurity compliance process, called the Authorization to Operate. The ATO process and its governing framework—the Risk Management Framework (RMF)—is ineffective, slow, duplicative, and in need of reform.[4]

A platform without resilient, adaptable software is increasingly brittle—and often irrelevant—in a fast-moving conflict. And software that is insecure is equally dangerous: Vulnerabilities in military systems are direct pathways for adversaries to degrade missions, manipulate data, or hold critical capabilities hostage. Recent conflicts make both points concrete. During the U.S. strikes against Iran, U.S. Central Command confirmed that operations extended explicitly into cyberspace.[5] And Russia's coordinated cyber and kinetic campaign against Ukraine's power grid induced an estimated five cumulative weeks of blackout.[6]

However, Ukraine's ability to adapt and continue fighting amid Russian cyberattacks did not stem from decade-long acquisition programs or cybersecurity compliance processes. Instead, it was enabled by rapid technology integration with support from some of the leading U.S. commercial technology firms—including Amazon, Apple, Cisco, Google, Microsoft, Palantir, SpaceX, and Tesla—demonstrating that speed, adaptability, and commercial integration are core elements in modern warfighting.[7] Conversely, bureaucratic delays can carry a devastating cost. A former intelligence officer describes just how high those stakes can be:

I firmly believed that software was the reason that a bunch of civilians had died. . . . We had critical fixes like known operational issues that were causing operational risk that were sitting on the shelf waiting to go through the ATO process.[8]

Efforts to modernize the ATO process began in earnest around 2013, when policymakers and practitioners recognized that it was not achieving the desired cybersecurity posture.[9] The ATO process yielded only static, point-in-time assessments instead of continuous risk evaluations. While initial reforms aimed to improve security outcomes, later efforts acknowledged that a slow approval process itself introduces security risks and delays the implementation of critical capabilities in the field. Despite a decade of reform—continuous authorization, DevSecOps platforms, Federal Risk and Authorization Management Program (FedRAMP) standardization—the ATO continues to be a major operational impediment in 2026.[10]

Methodology and Limitations

Prior research has examined elements of the ATO process in useful ways. Naval Postgraduate School researchers have analyzed RMF subjectivity, quantification challenges, and process inefficiencies within specific service contexts.[11] Government Accountability Office (GAO) and Inspector General audits have documented compliance failures, reciprocity gaps, and workforce shortfalls across individual programs and agencies.[12] This body of work has established important foundations.

This paper builds on that foundation while doing something distinct. It is the first analysis to examine not only the process itself but the foundational legal authorities, competing stakeholder incentives, and governance structures that collectively produce the delays the process reform literature has documented. This paper also assesses why reforms have not resolved those delays despite sustained attention over more than a decade.

To identify where bottlenecks occur, this research surveyed publicly available and government-internal repositories for baseline data on ATO timelines. No benchmarks were found—neither for overall approval timelines nor for time spent at individual stages of the process. In the absence of such data, this paper employs a qualitative, multisource methodology. It analyzes publicly available government documents, academic papers, congressional testimony, news reporting, and practitioner insights from key stakeholders in the process—including authorizing official (AOs), program managers, and system contractors—based on real-world experiences and use cases.

Available estimates of ATO timelines vary widely because practitioners, agencies, and reports use different assumptions, scopes, and start and end points. In the absence of consistent and transparent data, this paper can perform qualitative analysis, but it cannot perform a full value-stream analysis or precisely quantify the impact of individual reforms.

While the ATO process applies to all federal systems, this paper utilizes a realistic, hypothetical case study to illustrate the points of friction. The scenario follows a program manager seeking to field a commercial AI drone swarm capability. This capability is sourced via the Defense Innovation Unit (DIU). This organization was established specifically to transition commercial technologies to multiservice applicability, making it an ideal lens for this analysis. Because the proposed AI system also integrates raw intelligence data—rather than only finished intelligence analysis—it raises concerns for the intelligence agencies that collected the data, particularly around protecting sources and methods. Following this single capability through the

ATO process reveals how statutory authorities, delegated approvals, information opacity, and competing stakeholder incentives interact in practice, demonstrating why addressing any single bottleneck in isolation is insufficient. This same scenario is then carried through the proposed recommendations to illustrate how changes could alter the program manager's experience.

Background

While process-level pain points are documented in academic literature, congressional testimonies, and industry analyses, this section consolidates those friction points to expose easily overlooked foundational issues. Specifically, it examines how legal authorities, structural incentives, and opacity of information shape how the ATO process is executed. Understanding these underlying dynamics is essential to explaining why previous reform efforts have failed to streamline the process.

Pre-process

The Federal Information Security Modernization Act (FISMA) is the law that mandates a comprehensive information security program based on standards developed by the National Institute of Standards and Technology (NIST).[13] NIST developed the RMF, which provides the structured process agencies use to comply with FISMA. The process does not accredit technology in isolation. It evaluates the combined interaction of people, processes, and technology within a specific information technology (IT) system for a specific mission.[14]

A common misconception among program managers and commercial contractors is that a cloud Provisional ATO (P-ATO), an approved software list, or an existing ATO from another program constitutes blanket authorization for their mission through reciprocity. However, the way users interact with the system and data to support a specific mission is an integral part of the authorization. Consequently, deploying the same software for a different mission—involving different personnel and processes—currently requires a separate, dedicated ATO.

Responsibility for the ATO is also not borne solely by the government or by the contractors providing system integration or commercial software; it is distributed across multiple stakeholders. Figure 2 provides an overview of all RMF steps, associated actions, and major responsibilities, and it illustrates how “Step 0: Prepare” shapes every subsequent step. Step 0 establishes process requirements and expectations based on the organization’s policies and personnel. Failure to properly communicate these expectations in step 0 creates friction in every other part of the process.

Figure 2. ATO Process = NIST RMF Process

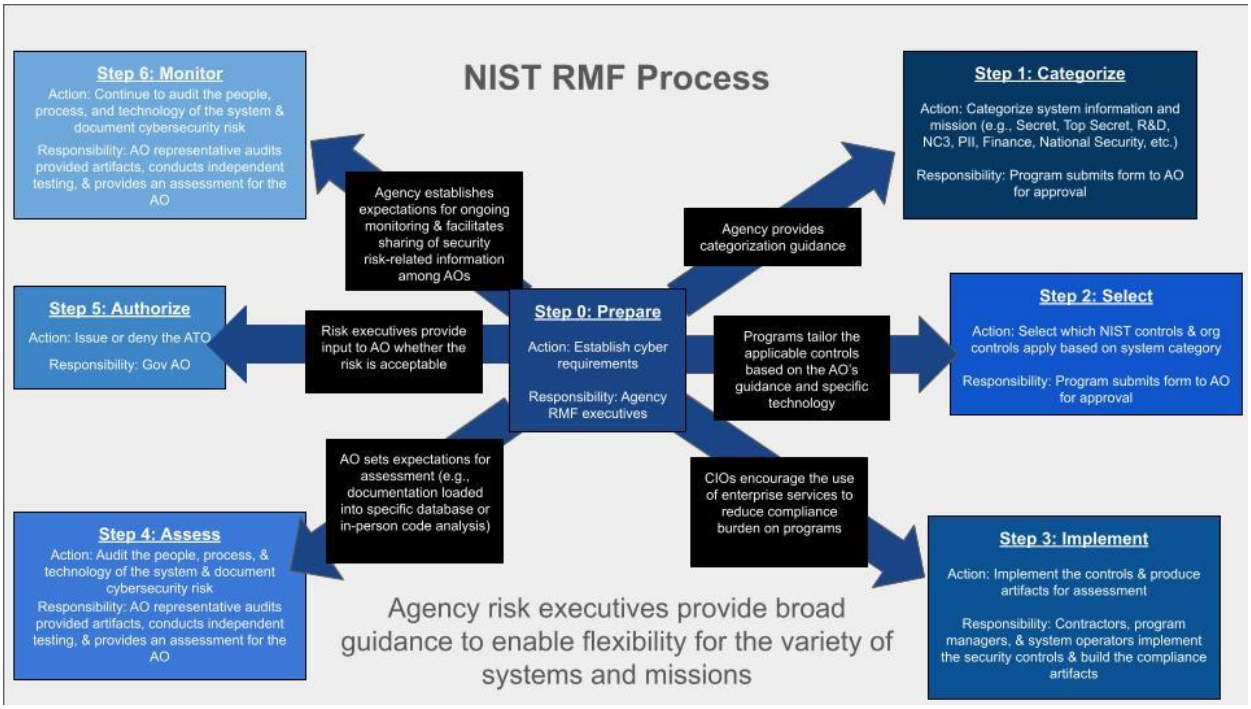

Source: Adapted from Computer Security Resource Center, "About the Risk Management Framework (RMF)," NIST, November 30, 2016, https://csrc.nist.gov/Projects/risk-management/about-rmf.

While the figure depicts a single organization working through one ATO, multiple organizations with multiple ATOs may be involved, making the process considerably more complex than this high-level view suggests.

When Multiple ATOs Are Required

FISMA applies to the entire government and assigns responsibility for protecting agency information systems and data to the head of the agency and the agency's chief information officer (CIO), with oversight from Congress and the Office of Management and Budget (OMB).[15] The U.S. Department of War (DoW)—also recognized at the U.S. Department of Defense—CIO delegates authority to the services and combatant commands, who further assign a specific AO based on mission or technology. An AO is the senior government leader who issues an ATO. Multiservice systems or systems expanding across fractured delegated authorities require independent or joint authorizations from all applicable AOs to maintain compliance.

A second trigger for multiple ATOs arises when a system hosts or processes IC Sensitive Compartmented Information (SCI). Even if a DoW AO has authorized the

system under their delegated authority, the presence of SCI data introduces IC equities that fall outside of the DoW's authority. In that case, an IC ATO is additionally required, governed by the Office of the Director of National Intelligence (ODNI) CIO, who delegates authority to different IC components. Figure 3 illustrates a small subset of these authority delegations and oversight relationships.

Figure 3. Example Delegated Authorities and Oversight

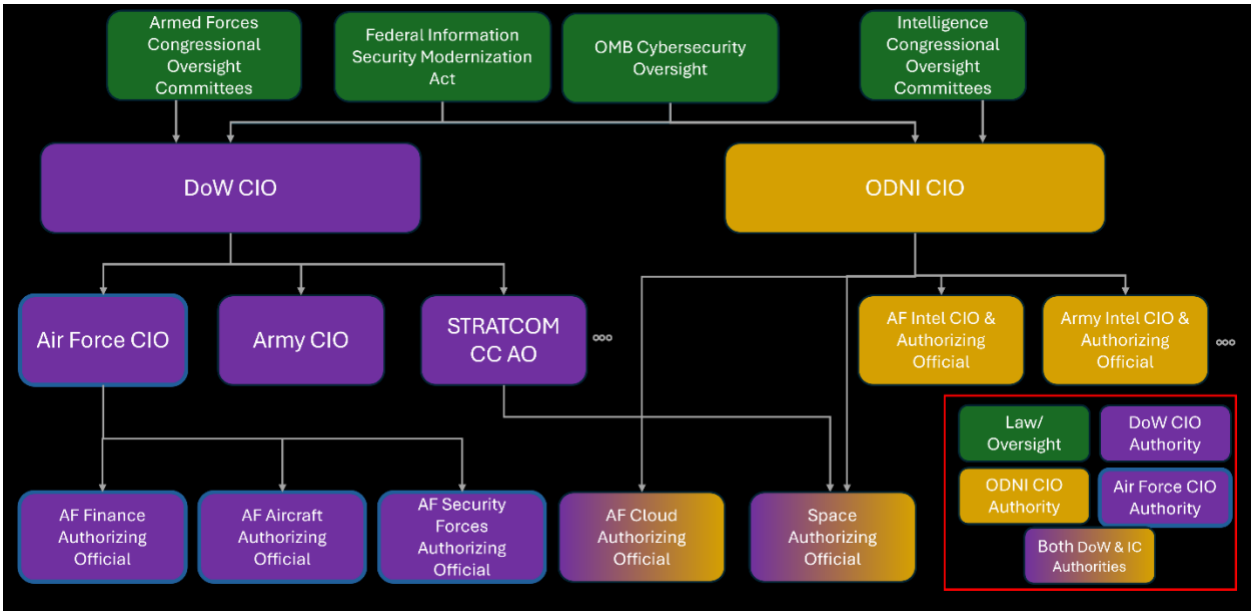

Source: Adapted from FISMA, U.S. Code Title 40 and Title 44, DoW Guidance.

In limited cases, some military AOs hold dual delegations from the DoW and IC CIOs. These AOs are accountable to both CIOs and are responsible for implementing the guidance and processes from each organization.

The separate statutory requirements placed on the DoW and IC CIOs can drive a single system or technology toward multiple ATOs. Because raw intelligence data is typically classified as SCI to protect sources and methods, its increasing integration into military operations will make dual military and IC ATOs a more frequent bottleneck to rapid system fielding.

But first, a program or company must identify an AO with appropriate authority and available capacity. This is the first point of friction in the process. Systems do not always fit cleanly under a single AO. For example, a system that monitors facilities on an Air Force base could be overseen by the Air Force civil engineering AO or the security forces AO or both, depending on who is employing the system. This complexity around delegated authorities creates additional bottlenecks.

Figure 4. Opaque AO Authorities and Scarcity

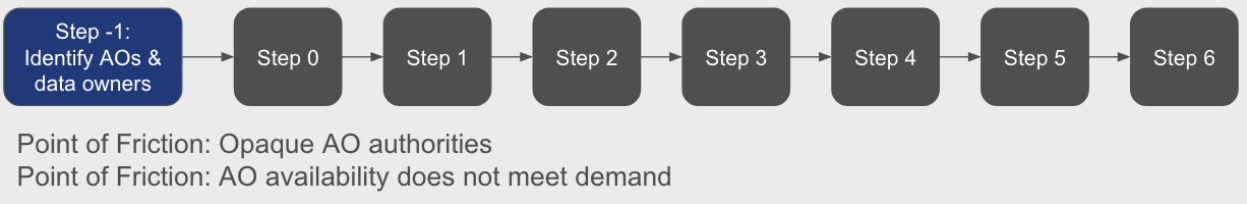

Source: CSET adaptation of Computer Security Resource Center, "About the Risk Management Framework."

AOs are a limited and overcommitted resource. Established programs with stable, multiyear funding and existing relationships are likely to receive higher prioritization from AOs. New commercial entrants—the nontraditional defense contractors that DIU was specifically created to attract—are unknown quantities to AOs and their staff, and therefore must compete for scarce review capability against programs that have been in the government ecosystem for years.

The pool of AOs and staff has not grown proportionally with the recruitment of new vendors. The FedRAMP—which is leading the federal government's automation and AI efforts through its FedRAMP 20x initiative—identifies its single greatest risk as the ability to hire and retain staff at scale, as shown in Figure 5.[16]

Figure 5. FedRAMP Government Staffing Over Time

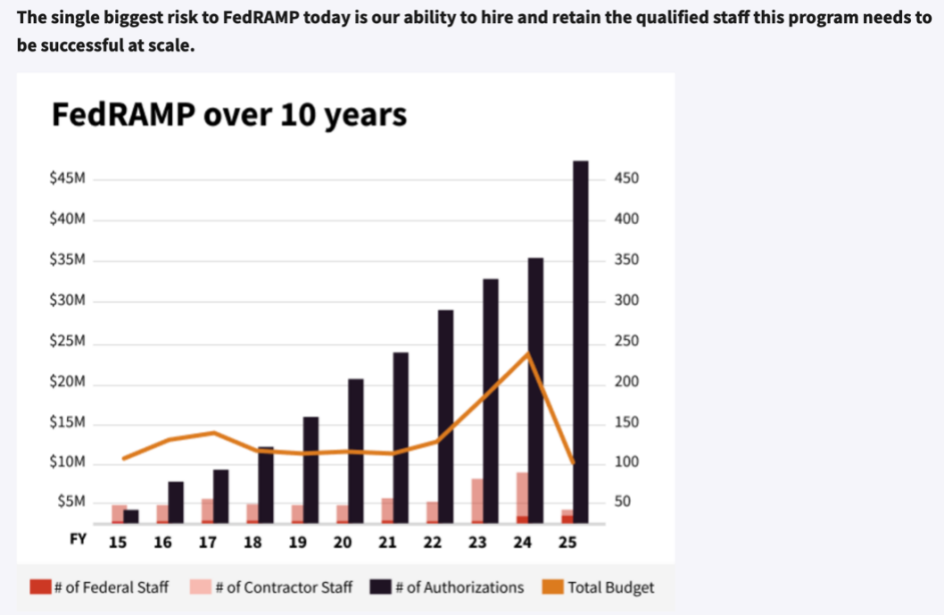

Source: FedRAMP, "FedRAMP Built a Modern Foundation in FY25 to Deliver Massive Improvements in FY26," news release, September 30, 2025, www.fedramp.gov/2025-09-30-fedramp-built-a-modern-foundation-in-fy25-to-deliver-massive-improvements-in-fy26/.

The GAO has documented this problem specifically for cloud services attempting to onboard through FedRAMP.[17] In the report, agencies have indicated that, although they wanted to use new cloud services, they were unwilling to take on the responsibility of sponsoring authorizations due to the associated costs, including hiring additional staff or contractors to facilitate the process. Because AOs lack the capacity to sponsor them, many commercial companies with secure cloud services cannot obtain FedRAMP authorization to work with the federal government.[18]

Independent third-party contractors can conduct cybersecurity risk assessments in step 4 of the process, but accepting risk and issuing an ATO under FISMA is a government function that cannot be delegated to contractors.

The staffing deficit creates a secondary problem known informally as "authorizing official shopping": programs and contractors searching for any AO willing to take them on. The absence of a centralized AO directory forces programs to rely on informal networks and inadvertently incentivizes seeking out the AOs who are perceived to demand less documentation, use less stringent risk assessments, and have the bandwidth to take on the work.

Vignette: Finding an AO

Imagine you are the DIU program manager for an AI-enabled commercial drone swarm capability that has Air Force, Army, and Navy units as transition customers. You know an ATO is required. Your first challenge is identifying at least one AO with delegated authority that covers at least one of your end users. Once you get an initial ATO, you hope that other AOs will accept the initial ATO through a process called reciprocity.

No consolidated, accessible directory of AOs and their delegated authorities exists. The operational units do not know who the appropriate AO is for their own mission, but you are able to curate a list through searching internal DoW websites and informal networking. Even after identifying officials with appropriate authorities, they may have a backlog stretching more than a year—as at least one DIU program has experienced firsthand. So now you must start shopping for your initial AO. Should you start with

- the Air Force aircraft AO because the system involves drones;
- the Army security forces AO because the system supports base defense;
- the Defense Information Systems Agency because of the AI-enabled command and control components deployable across services; or

- the Navy's research and development AO because the system has not yet been operationally accepted?

Since the system will be hosting SCI data, you will need to identify an AO who has delegated ODNI CIO authority, which is even more opaque. As an external program manager, you lack access to the secure networks and internal directories the IC uses to manage its authorization process, leaving you with no clear path forward. However, you hope that the IC AO will also leverage this initial ATO through reciprocity.

After months of aggressive networking and AO shopping, you successfully identify an Air Force AO with the appropriate legal authority for at least one of the intended missions and sufficient bandwidth to take on the ATO.

Competing Stakeholder Incentives

The military intentionally separates cybersecurity compliance authorities assigned to the DoW CIO from acquisition authorities assigned to the under secretary of war for acquisition and sustainment to ensure that security is not subordinate to cost and schedule pressures. That separation, however, sets the stage for no single actor to own the authorization problem. Program managers attribute fielding delays to AO review timelines; AOs cite incomplete documentation from contractors; contractors point to evolving ATO compliance requirements and access limitations.

The problem is compounded by a third stakeholder whose incentives are equally consequential and equally underappreciated: common control providers or enterprise IT services. The concept of shared IT services in the federal government predates the ATO process. The Clinger-Cohen Act of 1996 established agency CIO positions by law for the first time, directing federal agencies to focus on results achieved through IT investments and to promote an interoperable, secure, and shared government-wide IT infrastructure.[19]

Enterprise IT services implement and maintain shared security capabilities that multiple systems can leverage, including firewalls, identity and access management, logging, patch management, physical security, and network boundary protections. They also operate what are known as platforms: pre-secured computing environments that the government has already assessed and authorized, within which programs can build rather than certifying their own digital foundation from scratch. For software applications specifically designed to operate within one of these pre-secured platforms, enterprise services can automatically satisfy more than 80 percent of

required security controls, leaving the program responsible only for the smaller set of controls specific to its own application.

Government-owned platforms such as the Air Force's Platform One, the Navy's Black Pearl, and the Army's Enterprise Cloud Management Agency serve as enterprise services, as do select commercially operated platforms such as Second Front's Game Warden and Palantir's Foundry. These platforms themselves inherit controls from underlying cloud infrastructure that has already been assessed and authorized.

For unclassified cloud services, FedRAMP conducts a single standardized security assessment and publishes the results in a shared repository, allowing a cloud service assessed once to be used by hundreds of agencies without each repeating the full assessment. This practice is known as ATO reciprocity, formally defined as the "agreement among participating organizations to accept each other's security assessments, to reuse system resources, and/or to accept each other's assessed security posture to share information."[20] Table 2 provides a summary of the key stakeholders and their responsibilities, accountability, and incentives.

Table 2. Key Stakeholders' Incentives

Role	Responsibilities	Accountable	Incentives
Authorizing official	Accepting cybersecurity risk of a system and issuing an ATO	Agency CIO under FISMA	FISMA emphasizes preventing data breaches, not capability delivery. Structurally biased toward familiar technologies and risk avoidance.
Program manager	Balancing cost, schedule, and performance	Under Secretary of War for Acquisition and Sustainment	Evaluated on visible mission capability delivered on time and within budget. Cybersecurity compliance represents the absence of failure rather than a measurable achievement and can be expensive.
System contractors	Architecting and implementing system RMF controls and compliance artifacts	Variable based on contract terms	Generally incentivized to limit external dependencies and deliver contracted features. Expensive compliance improvements are less visible than new capabilities.
Enterprise IT services	Architecting and operating share security services across multiple systems	Variable based on contract terms	Incentive to standardize, stabilize, and reduce enterprise duplication. Generally prioritize established programs with stable, multiyear funding.

Authorizing Officials and Program Managers

The tension between AOs and program managers is perhaps the most visible friction point in the process. Program managers face constant pressure to quickly deliver operational capabilities to warfighters, and obtaining the ATO is almost always the primary bottleneck delaying deployment. Security and compliance requirements consume the limited program budget, and hardening a system to meet authorization standards can degrade the operational experience—particularly for commercial products not originally designed with government security constraints in mind. AOs, by contrast, bear no accountability for speed of delivery and carry no budgetary

responsibility. They can impose strict compliance requirements that consume a significant portion of a program's budget without absorbing any of the consequences when operational features are deprioritized as a result. Even when delays hold back critical capability, no mechanism exists to hold an AO accountable for the cost of that delay. As referenced in the introduction, a prior intelligence officer experienced a situation where civilian deaths may have been preventable with operational fixes that were sitting on the shelf, waiting to go through the ATO process.

Operational Users

Distribution of responsibility ensures that delays are easy to produce and difficult to attribute.

A less visible but equally consequential tension exists between the compliance framework and the operational user, who is not formally recognized in military instructions.[21] Unit personnel are far more likely to suffer from poor network performance, software failures, or compliance-driven configuration restrictions than from a targeted cyberattack. Their risk calculus is therefore fundamentally different from the framework designer's: They want a system that works reliably in their day-to-day mission, and compliance requirements that degrade performance are costs to minimize.

Figure 6 captures a representative sample of key authorization tasks and the stakeholders who share responsibility for them, illustrating that no stage of the process belongs to any single actor. Combined with the competing incentives described above, the distribution of responsibility ensures that delays are easy to produce and difficult to attribute.

Figure 6. Competing Stakeholder Incentives

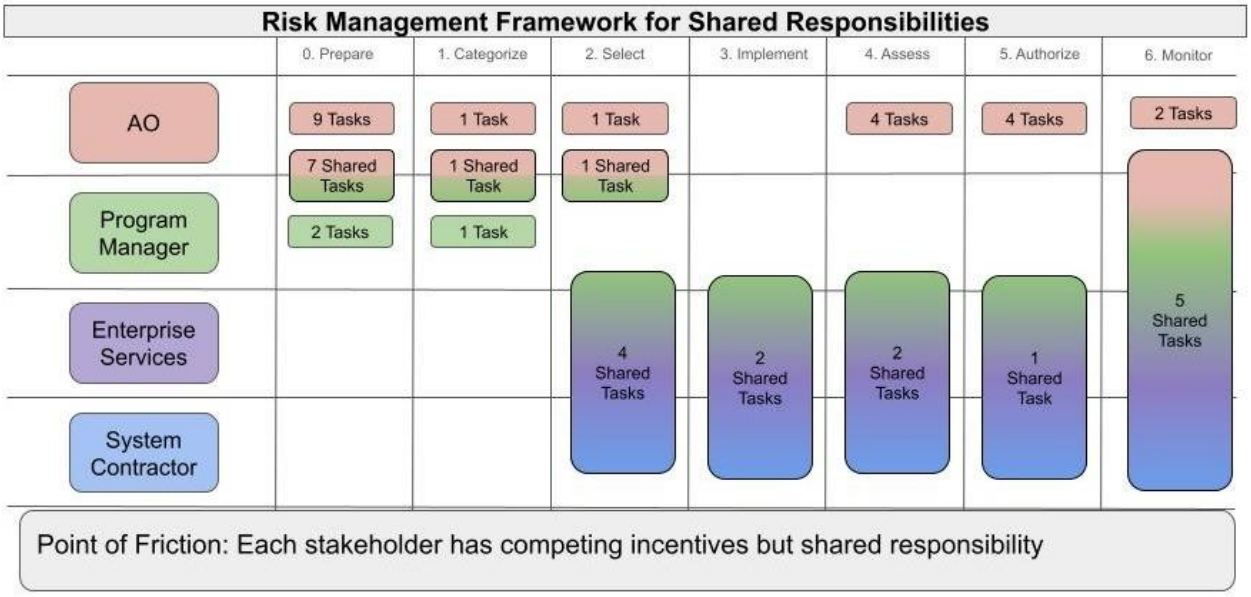

Sources: NIST, "RMF Roles and Responsibilities Crosswalk" (Department of Commerce, September 2024), https://csrc.nist.gov/csrc/media/Projects/risk-management/documents/Additional Resources/NIST RMF Roles and Responsibilities Crosswalk.pdf; Department of Defense CIO, *Risk Management Framework for DoD Systems*, DODI 8510.01, July 19, 2022, www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/851001p.pdf.

Vignette: Competing Interests

As the DIU program manager, you face intense pressure to field a minimum viable product well ahead of the DoW Software Acquisition Pathway's one-year delivery mandate, a timeline already compressed by the months spent merely identifying an available AO.[22] Accountable to the DoW CIO under FISMA, the AO naturally seeks to ensure robust security. The ATO is the AO's sole leverage to enforce cybersecurity because operational users often view security features as performance hindrances, while the commercial vendor strongly resists costly redesigns to their existing solution. Ultimately, the lack of standardized "minimum or critical controls" leaves the program manager, operational user, and vendor trapped in a subjective negotiation with the AO, trying to define a moving target for what is required to achieve the ATO.

Step 0: Prepare

Every organization must establish its own RMF guidance, polices, and requirements—and no two do it the same way. Departments, services, and individual AOs implement the RMF differently, with distinct templates and documentation formats and different interpretations of higher-level guidance, assessment methodologies, risk tolerances, and continuous monitoring expectations. The Department of the Air Force alone has maintained more than two dozen AOs aligned to different mission areas—including aircraft, finance, space operations, science and technology, logistics, civil engineering, manpower, operational test infrastructure, and enterprise IT—and they do not all have the same processes and paperwork requirements.[23] The Army and Navy maintain similarly segmented authorization structures across their mission portfolios.

NIST provides a comprehensive catalog of over 1,000 controls (that is, security requirements), explicitly intending for organizations to tailor this list to their specific mission rather than treating the entire catalog as a massive checklist. The problem is that AOs only care about a subset of controls that are rarely documented. Because these baseline expectations remain undefined, organizations are forced into a guessing game, resulting in profound inconsistencies in security postures and assessment expectations across similar missions. This unpredictability is compounded by severe inaccessibility. Even when AOs do establish tailored mission standards, process requirements, documentation formats, and assessment expectations, they are often inaccessible, incomplete, or overclassified. Some AOs host their required forms on internal portals that require service-specific credentials or URLs that demand separate permissions. The problem runs deeper than simple access: Unclassified ATO artifacts are sometimes managed exclusively within classified environments, and blank templates are marked as controlled unclassified information (CUI)—restrictions that bear no relationship to operational sensitivity.

The inaccessibility of AO process expectations—not cybersecurity requirements, which appropriately vary by mission and system—also creates upstream acquisition problems. Contracts cannot specify assessment expectations, documentation deliverables, or automation incentives when those expectations are unknown at award. A 2021 GAO review found that three out of five weapon-system programs reviewed did not include cybersecurity process requirements in their contracts at award, relying instead on vague stipulations to be "cyber resilient" or to "comply with the RMF."[24] Contractors told the GAO that such statements "do not provide enough information to determine what the government wants." The result is predictable: As AO expectations surface during the ATO process, contract modifications become necessary, schedules shift, and resources are diverted from mission features to newly discovered compliance

obligations. The AO's incentive structure—to avoid data breaches, not necessarily to adhere to delivery timelines—provides no check on this dynamic.

Programs often rely on informal networks simply to discover what is required of them. An internal survey by a cybersecurity firm that has executed RMF assessments across multiple federal organizations since 2016 found that timelines extended up to 36 months across 20 projects. More than half experienced credentialing or portal access delays that directly extended timelines. The projects that exceeded one year cited unclear or inaccessible AO expectations as a contributing factor, which leads to the next point of friction.

Figure 7. Incomplete, Inaccessible, and Non-standardized ATO Process Requirements

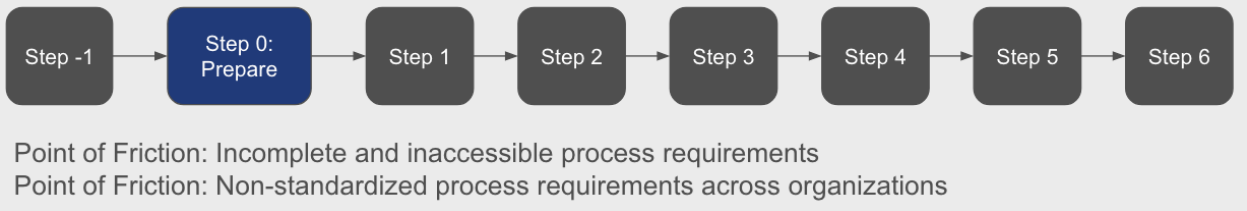

Source: CSET adaptation of Computer Security Resource Center, "About the Risk Management Framework."

Vignette: Incomplete, Inaccessible, and Non-Standardized ATO Process Requirements

With the right AOs identified, you now need their specific process requirements, and there is no official way to find them. You leverage your network until someone points you to separate documentation portals for each AO. You do not have the credentials to access either portal. After weeks—or months—of navigating two separate credentialing processes, you get in.

What you find is different from the last time you went through the ATO process for a previous program. One AO requires a specific set of static templates uploaded manually to a compliance repository, and another AO requires integration with their monitoring services and sends assessors to conduct direct code reviews. It is unclear if the assessors are familiar with the type of AI models used in the drone swarm software. With organizations using different templates, applying different assessment methodologies, and assigning different classification markings to

identical information, the future prospect of reciprocity is diminishing for the initial ATO.

Step 1: Categorize System

Step 1 in the RMF process is system categorization. The program must determine the sensitivity of the information the system will process, store, or transmit and assign a corresponding impact level that drives security requirements. This categorization decision involves interpretation, and different organizations, AOs, and program managers interpret the same guidance differently.

A concrete example illustrates the problem. Both impact level 4 and impact level 5 apply to CUI—sensitive but unclassified data requiring protection below the classified threshold. The distinction depends on whether the system qualifies as a "national security system," which triggers the impact level 5 designation and a more demanding and expensive set of security requirements. What constitutes a national security system is itself subject to interpretation. Some guidance allows a commander to designate their system a national security system based on perceived mission importance—and commanders naturally tend to view their systems as "critical to the direct fulfillment of military or intelligence missions."[25]

Figure 8. Varying System Categorization Interpretations

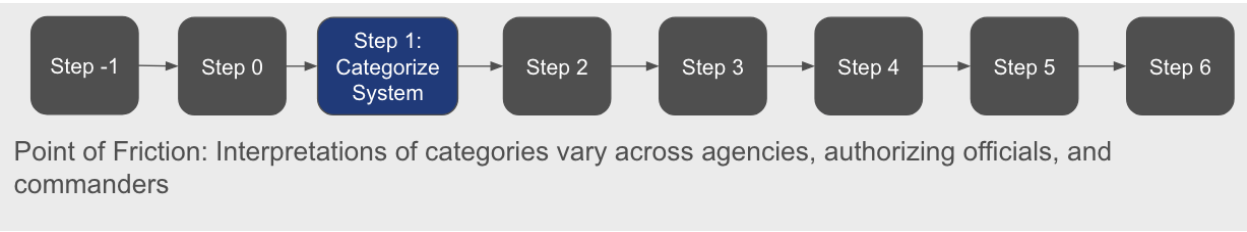

Source: CSET adaptation of Computer Security Resource Center, "About the Risk Management Framework."

Vignette: Varying System Categorization Interpretations

You need an unclassified development environment—a functional replica of the system with no classified data—so the contractor can develop and test updates before promoting code to the classified environment. That environment requires its own ATO. No clear guidance exists on whether a development environment constitutes a national security system. An unclassified development environment is

not operational and may not meet the threshold of being "critical to the direct fulfillment of military or intelligence missions" under NIST SP 800-59. But experienced program managers warn that upgrading from impact level 4 to impact level 5 mid-process is far more disruptive than starting higher, and most programs have already hardened their development environments to impact level 5 as a precaution, regardless of whether the designation is strictly warranted. You follow the same pattern, accepting a heavier compliance burden now to avoid a more costly restart later.

Step 2: Select Controls

Once a system has been categorized, the program must select the security controls it will implement—that is, the specific technical, operational, and management safeguards required to protect the system at its designated impact level. The full NIST control catalog contains over 1,000 potential controls.[26] For most military and IC systems, the applicable baseline runs to several hundred controls, each of which must be implemented, documented, and assessed before an ATO is issued.

However, control selection varies significantly among AOs and is rarely communicated transparently. This ambiguity makes the process feel like taking an exam without a syllabus, forcing programs and vendors to guess which controls the AO prioritizes. While some AOs demand comprehensive implementation up front, others accept a small set of critical controls. This systemic inconsistency makes predicting the minimum viable product required for an ATO highly challenging and dependent on the AO's personality.

Figure 9. Variation of Control Selection Across Orgs and All Controls Required Up Front

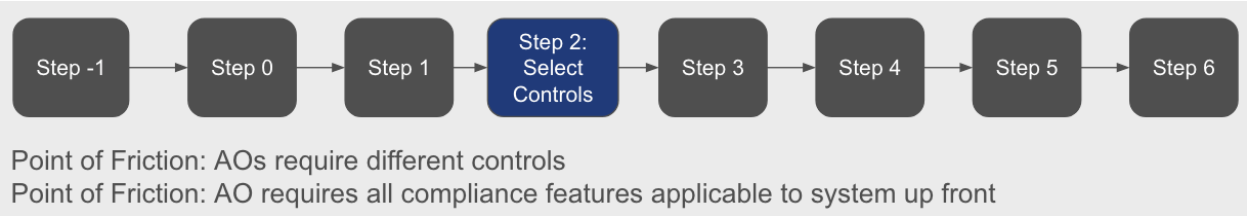

Source: CSET adaptation of Computer Security Resource Center, "About the Risk Management Framework."

Vignette: Variation of Control Selection Across Organizations

Once you and your contractor go through the thousands of potential security controls that could be applicable to the AI drone swarm, you are left with several hundred that you believe must be met to obtain an ATO. You submit the proposed control selection to the AO. The AO agrees that all these controls are applicable to your system and that you must meet all of them to get an ATO.

While you agree that all the controls are applicable, you try to get the AO and their staff to prioritize the controls and issue an ATO if you implement the most critical ones. Unfortunately, the AO wants everything implemented all at once because they are concerned that the program will not implement any further security measures after receiving an ATO. This clashes with how programs are supposed to build software: incrementally, releasing features to users in small batches and improving security continuously along the way (an approach called DevSecOps). Instead, the compliance team is stuck in an older, rigid mindset where every requirement must be fully finished before anything can move forward.

Step 3: Implement Controls

Implementation covers more than the deployment of technical tools such as firewalls, scanners, and encryption. It also includes writing policies for system access and outlining the operational procedures that personnel must follow on a regular basis to keep the system secure.

Once contractors establish the technical baseline and associated policies and processes, they must upload "artifacts" (evidence of compliance) into a repository for assessment. These artifacts range from automated scan results and access control policy documents to narrative descriptions of security processes. Automation tools exist to generate many of these artifacts, but contractors do not always implement them—sometimes due to lack of expertise, sometimes because software licenses are deemed too expensive.

When programs do attempt to use authorized cloud services, the programs frequently lack the permissions needed to automatically import the artifacts into their compliance packages, requiring manual transfers, sometimes coordinated through the cloud

service provider. Some enterprise IT services will support programs by uploading artifacts for reuse into the new compliance package; others do not. The technical capability to automate artifact transfer through application programming interfaces (APIs) exists, but government restrictions on repository access block the automation that would make this practical at scale.

Figure 10. Inaccessible and Manual Imports Without Automation

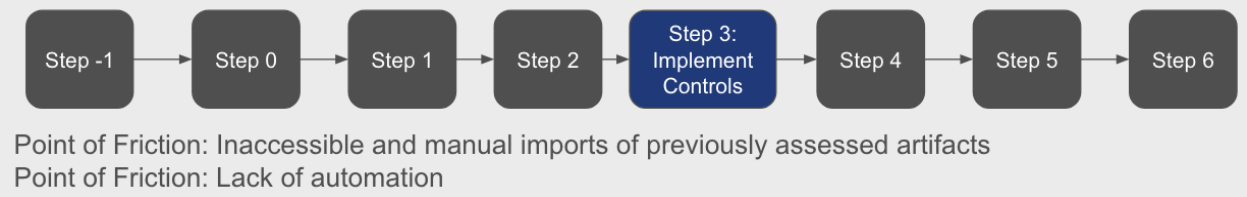

Source: CSET adaptation of Computer Security Resource Center, "About the Risk Management Framework."

Vignette: Implement Controls

Under time pressure—and knowing that the likelihood of the assessor being unfamiliar with the specific tools leveraged in your automated pipeline outputs is high—you and the contractor opt for static documentation. It is the path most assessors expect, so deviating risks confusion that could delay the review further. The package accurately reflects the system as configured at that moment but is not optimal for change management and may no longer be accurate by the time the package reaches the front of the AO's queue.

You also work on gathering the assessment reports for the authorized cloud services that you are inheriting as part of your infrastructure, so you obtain access with the appropriate CIO approvals and then manually upload the assessments into the new package. These assessments will also become obsolete and will not be automatically updated into your package when they get reassessed for new vulnerabilities and risks.

Step 4: Assess

Assessment methods vary widely, ranging from verbal authorization over the phone to paper-only reviews to physical system audits. Consequently, assessment quality depends heavily on the individual assessor and AO expectations. In the hands of technically competent and mission-aware assessors, this flexibility enables meaningful

risk evaluation. Without that expertise, it produces "compliance theater": completed checklists where security risks are neither identified nor understood within the mission context.

Context dictates risk: A technology that is insecure in one operational environment may be acceptable in another. Making that judgment requires a blend of threat knowledge, software expertise, and mission understanding. This hybrid skill set is rare in a single person and nearly impossible to staff consistently at scale across the military and IC. Highly skilled offensive cyber operations and vulnerability researchers are rarely RMF process specialists, and vice versa.

This critical skill gap is already evident with established technologies. For example, assessors routinely demand architecturally irrelevant physical hardware documents for cloud environments or struggle to comprehend containerization, a foundational DevSecOps technology.[27] AI-enabled systems will only deepen this divide. The highly specialized expertise required to meaningfully assess an AI model's training data, inference architecture, and decision boundaries bears no resemblance to the skills required to complete an RMF package. A workforce optimized for the RMF process, rather than for knowledge of modern security risks, is poorly positioned for success.

Figure 11. Static Assessment and Lack of Relevant Technical Knowledge

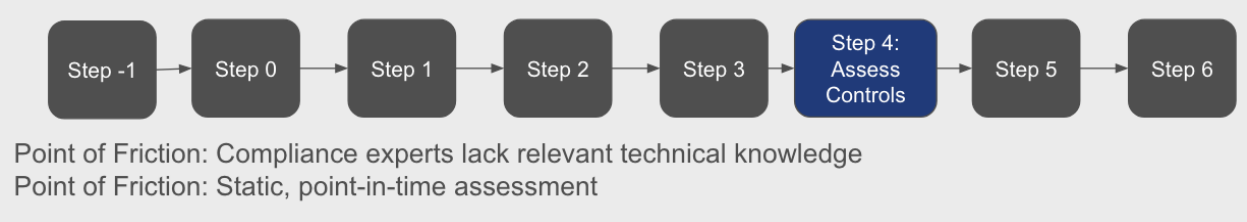

Source: CSET adaptation of Computer Security Resource Center, "About the Risk Management Framework."

Vignette: Static Assessment and Lack of Relevant Technical Knowledge

The AO's assessor arrives with one to two weeks allocated to review documentation, interview personnel, and verify that the system is configured as documented. The assessor is not familiar with the software architecture or with AI-enabled systems at all. What follows is not an assessment based on AI threats so much as a translation exercise: Engineers who are not compliance experts are attempting to communicate in an unfamiliar compliance language.

The AI model's security posture presents a particular challenge. Some of the information the assessor wants cannot be provided—not because it is being withheld, but because the nature of how the model works makes certain traditional compliance artifacts meaningless or impossible to generate. The two teams speak different languages, and bridging that gap requires weeks of back-and-forth, documentation rewrites, and improvised compliance translations. A month later, the assessment is complete. The package enters the AO's queue—and the waiting begins.

Step 5: Authorize

The AO's queue is typically long. AOs— frequently senior leaders at the general officer level or civilian equivalent—may be responsible for hundreds of systems, with only a handful of staff to support them.

AOs often lack deep technical familiarity with every system in their long and diverse queue, so the completeness and familiarity of the documentation become a proxy for the program's rigor and credibility. A package that "looks right" has an easier time than one that does not, regardless of the underlying security posture. Sometimes programs do submit incomplete documentation. A 2024 GAO report on FedRAMP noted that agencies reported cloud service providers frequently lacked complete documentation.[28] Whether that reflects contractor negligence, a process too complex to navigate without insider knowledge, or expectations never clearly communicated in step 0 is difficult to untangle. It is often some combination of all three.

The structural incentives of the AOs will also compound the delay. Saying no, requesting additional documentation, or deferring a decision is legally defensible and institutionally safe, whereas accelerating authorization can put the AO at risk in case of a breach. While AOs can issue ATOs that are valid for up to three years, few programs can meet all of the controls that help an AO feel comfortable with accepting the risk for that long. Moreover, if the program is not funded past a year, the ATO may only be valid for that long, given that the program will not have the funds to sustain a healthy security posture.

Figure 12. Long AO Queues and Incomplete Documentation

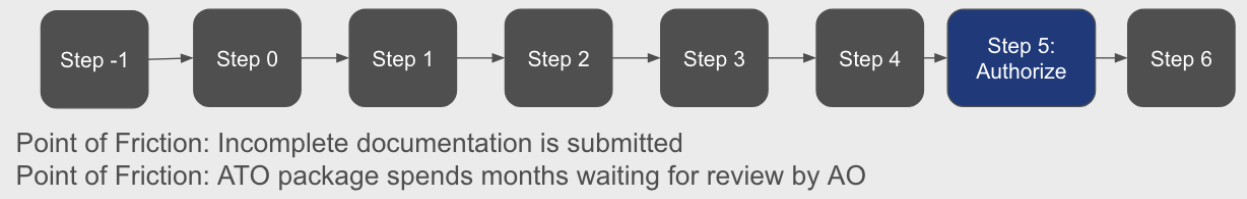

Source: CSET adaptation of Computer Security Resource Center, "About the Risk Management Framework."

Vignette: Long AO Queues and Incomplete Documentation

After the assessment concludes, your package enters the AO's queue. Months pass. When it is finally reviewed, the package is returned—not because the system is insecure, but because documentation the AO expects to see is missing. The expectation was never clearly communicated in step 0, or was buried in a portal you couldn't access, or was simply overlooked in the rush to submit ahead of your planned fielding date to satisfy the AO's lead-time requirements. The package is revised and resubmitted, and the wait begins again.

Eventually you receive an ATO. It expires in a year because your funding is not guaranteed beyond one year. The AO views your ability to sustain a heathy security posture beyond the funding guarantee as a risk.

Step 6: Monitor

Receiving an ATO is not the end of the process but rather the beginning of an obligation. Step 6 requires continuous monitoring of every applicable control for the life of the system: tracking configuration changes, assessing control effectiveness, documenting deviations, and reporting the system's security posture to designated officials. In theory, this produces a living security picture that keeps pace with an evolving threat environment.

Continuous monitoring became a formal requirement in 2014. Over a decade later, it remains one of the least consistently executed elements of the RMF. Monitoring can be manual and labor-intensive, and it competes for the same limited AO staff already managing hundreds of systems through the authorization process. For programs, sustaining a rigorous monitoring posture requires dedicated compliance auditors and automation that was not built up front. Furthermore, continuous monitoring comes

with incident response requirements. However, contention arises when security operations centers—the teams that monitor systems for threats—request administrative access to mission applications they do not fully understand. Granting that level of access can backfire: without deep knowledge of how the application works for a given mission, a security operations center could accidently cause the very outage it is supposed to help prevent. The result is that continuous monitoring is rarely achieved, and AOs default to one-to-three-year ATOs and revisit systems only when the clock expires.

The problem is not invisible to leadership. Then acting DoW CIO Katie Arrington put it plainly: "An ATO is granted at a very specific time in the network, the architecture of the network, the iteration of the software. Everything is like a snapshot in time, it's a static moment. But software is dynamic, it changes—every patch, every iteration, every version. So why wouldn't we move to a continuous ATO?"[29] In 2025, she felt compelled to declare she was "blowing up the RMF, blowing up the ATOs. They're stupid. They're archaic."[30] This statement, more than a decade after continuous monitoring was mandated, is perhaps the clearest measure of how far execution has fallen short of intent.

Figure 13. Manual, Periodic Monitoring and AO Scarcity

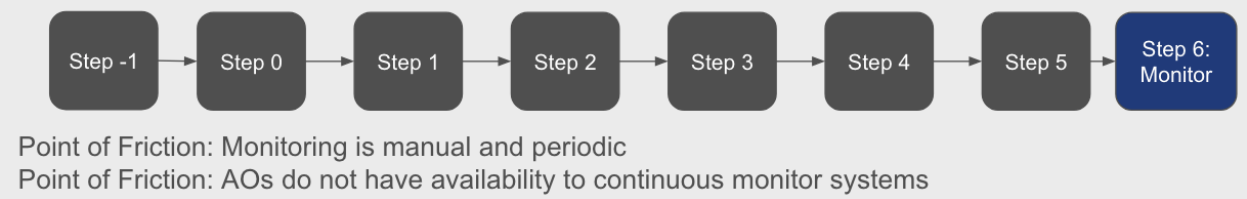

Source: CSET adaptation of Computer Security Resource Center, "About the Risk Management Framework."

Vignette: Manual, Periodic Monitoring

Shortly after receiving the ATO, the system requires software updates: new features, patches, and architectural changes driven by operational feedback from the field. Under the current requirements, significant changes must be documented and potentially resubmitted to the AO for review. But "significant" is not defined with precision, and your AO's staff is already stretched thin and has no bandwidth for a reassessment.

Due to the lack of automation, both you and your AO lack real-time insights into the risks these changes introduce. You face a choice with no good answer. You

can push the updates and ask for forgiveness later, relying on a subjective interpretation of what constitutes a significant change, or hold the updates in the queue and wait months for the AO's review, leaving operators in the field with a system that does not incorporate the fixes they need right now.

Reciprocity

An ATO issued by an AO within in one military service does not automatically extend to another, due to the delegated authority structure illustrated in Figure 3. Reciprocity is intended to reduce administrative burdens for a second ATO package by allowing the reuse of assessment results and documentation from the initial ATO, but three main obstacles cause reciprocity to fail.

First, AOs require different controls and documentation across the military, even for similar missions and technologies. This means a program can build the same capability for two different military branches, or even two units within the same branch, and still have to satisfy separate, often non-reusable sets of paperwork and security requirements for each one, even though the underlying technology and risks are similar, if not identical.

Second, variations in systems and their employment make it difficult to assess the original ATO and determine if it is appropriate for reciprocity. To start, discoverability is limited: Programs and AOs cannot search across system compliance repositories, making it difficult to identify existing ATOs available for reuse. Even when a potential ATO is identified, obtaining the packages requires access that must be manually approved by the original AO. Once access is granted, the receiving program and AO must manually parse the package to identify relevant controls. This task is especially difficult when controls involve shared responsibility, a common challenge with cloud and enterprise IT services. Then the AO must evaluate the previous assessment's findings against their own organization's risk tolerance.

Third, there is no strong incentive for AOs to leverage reciprocity. The 2024 DoW *Cybersecurity Reciprocity Playbook* encourages reciprocity by directing AOs to escalate reciprocity assessment disputes to the DoW CIO. Nevertheless, ATO packages are so tightly access controlled that AOs, program managers, and vendors are sometimes unaware that a reusable assessment exists.[31] Consequently, program offices frequently abandon reciprocity, start new packages from scratch, and manually re-upload prior documentation. Reciprocity becomes replication with extra steps.

Figure 14. Reciprocity Variation, Unassed Risk Deltas, and Lack of Incentives

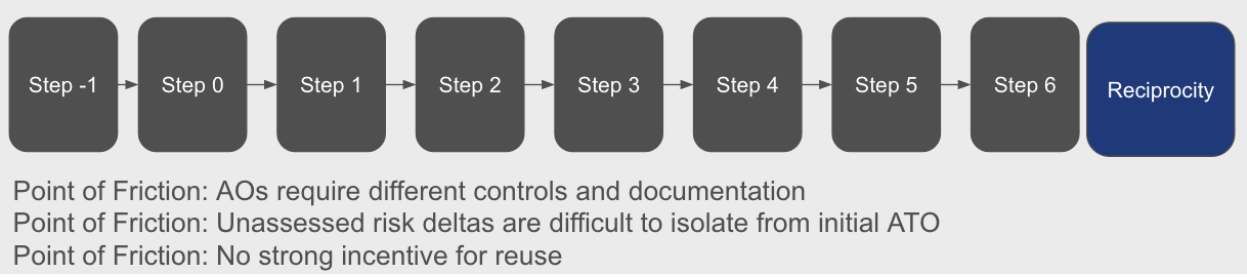

Source: Source: CSET adaptation of Computer Security Resource Center, "About the Risk Management Framework."

Vignette: Failed Reciprocity

Now you need to obtain the IC ATO for the SCI data, and you are working to field the system with an Army unit and hope to leverage the Air Force ATO. Fortunately, the existence of the initial package is known. Furthermore, because the Army security forces unit intends to employ the system in a similar manner, the original ATO package is largely applicable. The Army security forces AO agrees to review the existing package. However, upon obtaining access, the Army AO struggles with the Air Force's formatting, values different controls, and prefers a different assessment methodology. Additionally, it remains unclear which parts of the package must be modified to support the specific unit.

Consequently, the Army AO disagrees with the initial assessment and declines to issue an ATO. This prevents the Army unit from fielding the system for operational use until the dispute is resolved through the DoW CIO escalation process. Because operational demand in the field is critical, instead of waiting for resolution, you build a new ATO packaged tailored to the Army AO's preferences, manually incorporating the Army-specific modifications.

The IC ATO process also values different security control requirements, documentation standards, and assessment methodologies. The number of authorization packages required is now expanding in direct proportion to the program's operational success. The authorization process is a constant constraint, and it scales poorly precisely when the commercial technology scales across the military well.

Reform Analysis

The ATO process has not lacked for reform attempts. Congressional interest, military innovation, and federal program modernization have produced a decade of initiatives, each addressing real problems and each falling short of resolving them. This section evaluates those reforms against the friction points documented in the Background section, identifies where reforms have made meaningful progress, and surfaces the gaps that current efforts leave unaddressed.

Bottleneck Analysis

As mentioned in the Methodology and Limitations section, there is no consistent, government-wide measure of how long the ATO process takes—and that is itself a problem. Reported timelines vary widely depending on the background and incentives of the reporters. Program managers start the clock from contract award. AOs count from when a complete package lands in their queue. Contractors count from when they receive clear compliance requirements, which may be long after the contract begins. Reform advocates cite cases of rapid ATOs for programs with mature DevSecOps pipelines that build upon a previous ATO. Critics cite multiyear timelines for first-time authorizations with novel commercial technologies and architectures. All of these accounts may be accurate within the scope of what each stakeholder controls and observes.

That framing obscures where the real delays originate. As documented in the Background section and corroborated by practitioners at multiple stages of the process, programs frequently spend months before step 0 identifying an AO with appropriate legal authority, obtaining portal credentials, and discovering process requirements through informal networks.

Figure 15 illustrates a notional distribution of time across the process based on reported experiences from practitioners at various stages: program managers, contractors, assessors, and AO staff. The various experiences are illustrative rather than statistically definitive, and actual timelines vary significantly based on mission type, classification level, the AO's familiarity with the program, and whether the program is pursuing an initial ATO or reauthorization. These caveats matter, and the figure should be read as an indicator of where the biggest bottlenecks may be.

Figure 15. RMF Timeline Analysis

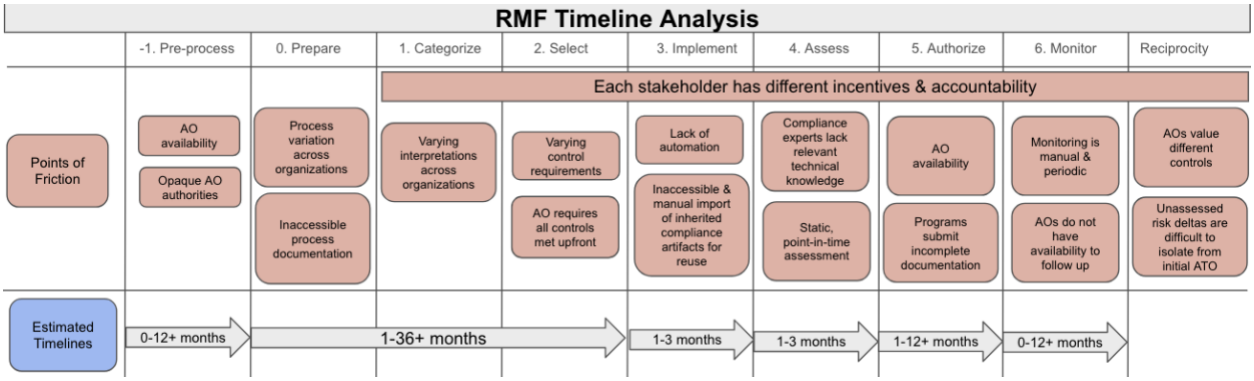

Source: Author's analysis based on 2019 *GAO Cloud Computing Security;* 2024 *GAO Cloud Security;* 2025 *FedRAMP Built a Modern Foundation* and practitioner experiences.

The first half of the ATO process accounts for a disproportionate share of total elapsed time

With those caveats in place, the pattern is consistent enough to be analytically useful. The first half of the ATO process accounts for a disproportionate share of total elapsed time for programs pursuing a first ATO. That finding indicates that steps -1 through 2 are the highest-leverage intervention point for a first-time ATO for a new commercial company like the AI-enabled drone swarm.

However, the calculus changes substantially once a program has already achieved an ATO and established the people, processes, and automation required to sustain it. At that point, the bottleneck shifts. The pre-process friction is largely resolved: The AO relationship exists, the portal credentials are in hand, and the documentation expectations are understood. The time savings from automation, continuous monitoring, and DevSecOps pipelines become the dominant variables after the ATO is granted. This context is essential for interpreting the reform landscape accurately. Nearly every claim of dramatic ATO acceleration—"ATO in a day," "90-day authorization," "continuous ATO"—refers to programs that have already made the up-front investment in DevSecOps infrastructure, established AO relationships, and built the compliance automation that makes rapid reauthorization possible.[32] These are meaningful achievements, but they do not account for all of the friction points that exist for a company that is not an established enterprise IT service and requires a first-time authorization.

The implication for reform prioritization is significant: Reforms that accelerate steps 3 through 5 deliver high returns for programs already inside the system. The highest-return reforms for first-time authorizations are those that address the first half of the process.

Reform Descriptions

Congressional actions and military initiatives have addressed overlapping but distinct aspects of the process. Understanding what each reform targets—and what it does not—is necessary context for the gap analysis that follows.

Congressional Reforms

Congressional concern about the ATO process has built incrementally. The fiscal year (FY) 2021 National Defense Authorization Act (NDAA) required the military to insert software security criteria into acquisition solicitations.[33] This is a step toward holding contractors accountable for their role in the process and clarifying expectations before a contract begins. The intent was sound, but as the GAO has documented, requirements that simply direct contractors to "comply with RMF" provide insufficient specificity when the process itself is opaque and AO expectations are inaccessible.[34] The reform addressed an accountability gap without closing the underlying transparency gap that makes accountability difficult to enforce.

The FY2023 NDAA marked the first significant legislative action on reciprocity, codifying FedRAMP into law and establishing a presumption of adequacy—meaning a cloud provider authorized by one agency could have that authorization accepted by others.[35] The provision applied only to unclassified cloud environments, limiting its applicability to military systems operating at higher classification levels.

The FY2025 NDAA went further, directing the military to develop policy requiring officials to accept security analysis and artifacts already authorized by another military component, and to standardize accreditation documentation to enhance reciprocity, with the explicit goal of allowing more rapid cloud capability adoption without redundant authorization.[36]

The FY2026 NDAA represents the most direct congressional intervention to date. Section 1521—titled "Accountability of the Authorization to Operate Processes"—established department-wide mandatory timelines, expedited review and appeal procedures for delayed authorizations, and required escalation processes when approvals stall, written justification for delays, and biannual reporting to Congress on ATO decisions, denials, workforce qualifications, and recommendations to improve the

RMF.[37] The significance of this goes beyond process efficiency: It is the first time that the FISMA framework's emphasis on confidentiality has been explicitly balanced against speed. By imposing accountability for delays, Congress has begun to rebalance the asymmetric incentive structure that has structurally biased AOs toward risk avoidance. Whether the timelines and escalation mechanisms will be implemented with sufficient rigor to change behavior remains to be seen.

It is worth noting that congressional reforms have focused heavily on cloud services receiving P-ATOs and inherited into systems. Programs that do not follow the cloud-native architecture that these reforms assume—which may include AI-enabled systems with edge hardware not in the cloud—fall largely outside the scope of what these reforms address.

FedRAMP Reforms

FedRAMP 20X represents the federal government's most ambitious attempt to automate the cloud authorization process, with the General Services Administration (GSA) pursuing machine-readable security documentation, automated artifact validation, and AI-assisted assessment to dramatically reduce the manual burden that has made FedRAMP underutilized despite its mandate.[38] However, FedRAMP 20X applies to unclassified cloud services and does not extend to the classified environments where most military operational systems reside.

Military Reforms

The concept of continuous authorization hit the news in 2018, when the Air Force software development unit Kessel Run established a continuous ATO to pursue a fundamentally different approach to updating and releasing new software: one that authorized the development pipeline rather than static system snapshots, allowing developers to continuously push validated updates without triggering a new authorization for every single update.[39] The fast-track ATO followed in 2019, oriented toward applications deploying to secured cloud infrastructure through enterprise IT service inheritance. This accelerated the authorization of cloud-native applications that could leverage inherited controls from an enterprise IT service provider for a significant portion of their baseline.[40]

The 2024 Continuous Authorization to Operate (cATO) evaluation criteria formalized continuous authorization requirements for the entire DoW, specifying that qualifying systems must demonstrate continuous monitoring of security controls, active cyber defense measures, and adoption of DevSecOps practices.[41] Critically, these reforms did

not make the initial ATO easier to obtain. They made the process of pushing updates significantly faster after an initial ATO is achieved—a meaningful improvement for programs already inside the system but offering limited relief to programs still navigating the pre-authorization friction documented in the Background section. The cATO criteria do, however, create incentives for automation and continuous monitoring that address several downstream pain points.

Most recently, in September 2025, the DoW released the Cybersecurity Risk Management Construct (CSRMC).[42] Designed to shift from static compliance checklists and manual processes to continuous monitoring, automation, and cyber defense at operationally relevant speed, the CSRMC is grounded in ten core principles: automation, critical controls, continuous monitoring, DevSecOps, cyber survivability, training, enterprise services and inheritance, operationalization, reciprocity, and cybersecurity assessments. Figure 16 illustrates how the CSRMC maps to RMF steps. The CSRMC is the most comprehensive reform to date in terms of the number of friction points it addresses, but it was released recently enough that implementation effectiveness cannot yet be assessed.

Figure 16. Cybersecurity Risk Management Construct

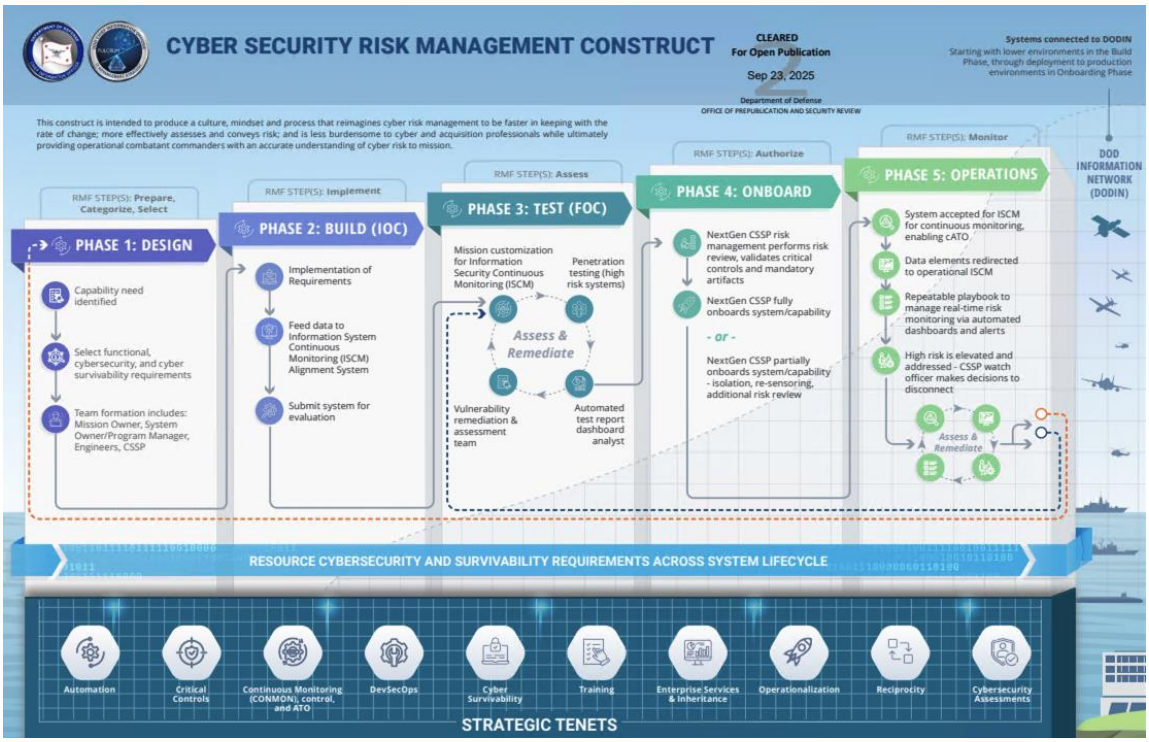

Source: Department of War, "Department of War Announces New Cybersecurity Risk Management Construct," news release, September 24, 2025, www.war.gov/News/Releases/Release/Article/4314411/department-of-war-announces-new-cybersecurity-risk-management-construct/.

This reform does not appear to address a critical need to reduce friction for new commercial software that could prove essential for the military. The CSRMC's potential appears strongest for a specific class of systems: cloud-hosted, containerized applications that fit cleanly within established enterprise IT services. For commercial software that was not originally designed to meet government security standards, the path to compliance remains architecturally demanding regardless of the framework.

A more nuanced challenge exists on the other end of the commercial spectrum. Some commercial companies—particularly those in high-threat industries like financial services and critical infrastructure—operate sophisticated, continuously evolving security architectures precisely because they are under constant attack and financially incentivized to stay ahead of adversaries. The government does not always have access to the most current commercial cybersecurity tools and infrastructure. Ironically, this is sometimes the case because the authorization process required to adopt them takes longer than the threat environment that motivated their development. As one Amazon Web Services employee observed: "When AWS constructs the same services across various partitions—including commercial, GovCloud, secret, and top secret—the lengthy ATO process can actually hinder our ability to provide more secure architectures in other partitions, simply because it takes so much time to comply with all the necessary standards." The process intended to grant access to commercial security innovation can, in practice, block it.

Systems that operate at the tactical edge and do not appear to fit within the current enterprise IT services—such as the AI-enabled drone swarms in the vignettes—present a further challenge. Whether the CSRMC's benefits will extend meaningfully to cutting-edge AI systems remains an open question. While the reforms have good intentions, the resulting lack of real observed change comes from the deeper process-oriented facets of RMF implementation across DoW and the inherent nature of a disaggregated, large scale, non-standard enterprise to navigate.

Connecting Reforms to Friction Points

Figure 17 maps each reform to the steps of the RMF process that it is focused on. The results of this mapping then influenced this paper's recommendations.

Figure 17. Reform Mapping

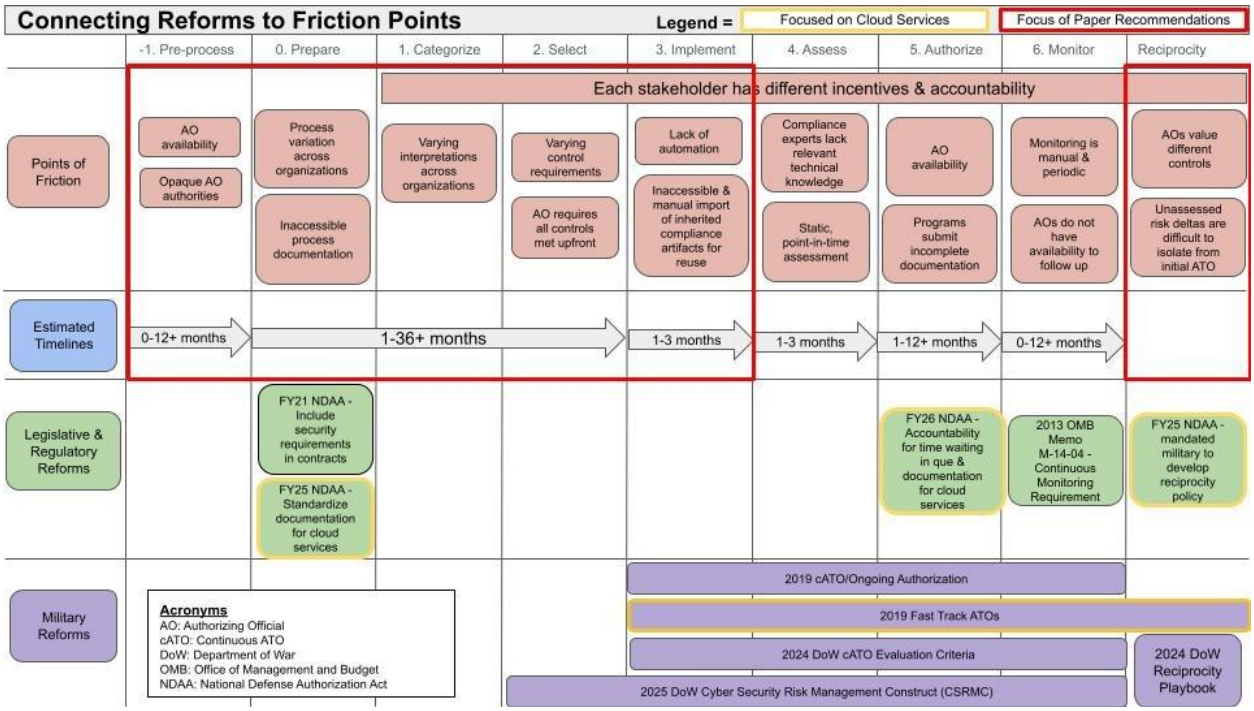

Source: Author's analysis based on 2013 OMB Memo; 2024 GAO *Critical Cybersecurity Challenges;* 2019 *GAO Cloud Computing Security;* 2024 *GAO Cloud Security;* 2024 *DoW Cybersecurity Reciprocity Playbook;* 2021 *GAO Weapon Systems Cybersecurity;* 2016 NIST *About RMF;* 2021, 2023, 2025 and 2026 NDAAs; 2024 cATO Evaluation Criteria; 2019 AF *Fast-Track ATO;* 2025 *FedRAMP Built a Modern Foundation;* 2025 *DoW Cybersecurity Risk Management Construct* and practitioner experiences.

Current reforms have left the first half of the process almost entirely unaddressed, which is significant, because that is where new commercial entrants face the greatest burden. Most reforms also address process execution rather than structural incentives, which limits how much lasting change they can produce. The FY2026 NDAA is the only reform that directly modifies the accountability structure by imposing consequences for delay and requiring written justification for denials. Every other initiative attempts to make the existing process faster or more automated without changing the incentive environment.

While the DoW *Reciprocity Playbook* references reusing IC artifacts for military authorizations, process reforms on the DoW side have no impact on the IC's authorization requirements. FISMA's vertical accountability structure assigns responsibility within agencies, not across them, and the FY2026 NDAA's requirements apply to the military but not to the IC AOs accrediting the classified services that military programs also depend on. As AI systems increasingly require fused intelligence and operational data, dual-authorization scenarios will become more

common rather than less. This will require collaboration between the DoW and IC CIOs—or, more likely, OMB or congressional intervention to establish cross-agency accountability mechanisms that FISMA's current structure does not provide.

The IC's processes can be faster due to flatter organizations structures, one AO organization to navigate, a single process, and mandated enterprise IT services with a stable source of security policy and funding. But a lack of standardization between the IC and DoW regarding the ATO process leads to duplication and sometimes disjointed security postures.

The mapping in Figure 17 reflects the friction points CSRMC is designed to address based on its stated principles. CSRMC may have not-yet-public implementation plans that address additional pain points, and its actual impact on authorization timelines will only become clear as adoption matures across the services.

The following section identifies the highest-return opportunities available given these observations. These changes are relatively achievable, complement rather than duplicate existing CSRMC and NDAA efforts, and target the friction points current reforms have left unaddressed. These are not exhaustive recommendations for every single point of friction, but a prioritized set of interventions designed to produce the greatest impact given current AI capabilities and policy reform momentum.

Recommendations

The AI-enabled drone swarm vignette that ran through the Background section returns here. Where it previously illustrated what the process costs—in time, resources, and operational risk—it now illustrates what a reformed process could look like in practice. The vignette is not a guarantee of what the recommendations would deliver. It is an illustration of what becomes possible.

Recommendation 1: Publicly Release Critical Controls and Authorizing Officials

Federal organizations should publish their AO-delegated authorities and critical controls, and they should help programs and vendors identify an appropriate AO with bandwidth. This does not require revealing the identity of individuals, system vulnerabilities, risk determinations, or specific implementation details.

Following principle in cryptography—where algorithms are public, but keys are kept secret—publishing ATO critical controls does not expose vulnerabilities or their specific implementation. Instead, it establishes clear, public standards against which and evaluated. Publishing delegated authorities would allow for AI to assist with AO identification. Opacity creates administrative drag without enhancing security.

Clear minimum standards provide predictability for industry. Traditional defense contractors and nontraditional commercial companies can assess compliance expectations before committing resources to development, enabling more accurate proposal pricing and more realistic schedule planning. Conversely, administrative barriers discourage the commercial innovation the military needs. In contrast, China's military-civil fusion strategy is currently dramatically expanding the pool of non-state companies supplying the People's Liberation Army.[43]

The CSRMC is already moving in this direction by identifying critical controls. This signals a recognition that not all controls carry equal security value and that the DoW must prioritize its efforts. The next step would be to publicly release these critical controls along with AO-delegated authorities and process requirements. Making this data discoverable will eliminate the months that program managers and vendors currently spend navigating administrative opacity.

Vignette: Recommendation 1

With this reform in place, the authorization process begins before the contract is written, rather than months after it is awarded. You locate the published delegated authorities for both the military and IC AOs relevant to your drone swarm capability. Even though you still have to approach each one to inquire about their availability, the search is faster and more comprehensive.

You put the published critical controls directly into the contract solicitation for proposals, replacing the generic "comply with RMF" language that previously gave contractors no actionable guidance and gave you no accountability lever when requirements emerged mid-process. Even better, the contractor arrives already knowing what is expected, with a cost and schedule that reflects actual authorization requirements rather than a placeholder carrying undisclosed risk.

Recommendation 2: Leverage AI to Standardize to Machine-Readable Formats

Reducing duplication, increasing automation, and enabling genuine reciprocity requires more than encouraging AOs to share their critical controls and authorities; it requires standardizing the format in which those requirements are expressed so that they can be automated. AOs should be required to feed their process standards and documentation templates into a centrally managed, AI-enabled tool designed to normalize ATO submissions.

A well-designed AI tool or agent accessing all the AOs' process requirements and documentation would accomplish several things simultaneously. It would normalize RMF documentation forms and control taxonomies across AOs, assist program managers and contractors in generating contract language, and enable the automated population of standardized forms based on system architecture data.

More fundamentally, the military and IC should adopt Open Security Controls Assessment Language, the machine-readable format developed by NIST. Using OSCAL would shift compliance away from static Word documents and PDFs toward structured, interoperable data.[44] OSCAL-formatted packages can be read, validated, and transferred by automated tools without manual repackaging, dramatically reducing the effort required to prepare documentation for multiple reviewing authorities. AI tools can translate OSCAL outputs into human-readable formats for AO staff, bridging the gap between automation and human review without requiring

compliance personnel to read structured data. This makes machine-readable compliance the default output rather than an afterthought requiring additional conversion work.

An alternative approach would allow each AO to maintain unique processes and rely on AI to translate between them—similar to how tax software translates between different state and federal tax codes. While technically feasible, this trades one problem for another: Translation accuracy degrades as process variation increases, and AI-generated translations of ambiguous or conflicting requirements introduce risk rather than reduce it. Standardization with a machine-readable language at the source is a more durable solution than translation after the fact.

The compounding benefit of simplifying compliance documentation is often understated. Every hour a skilled engineer spends reformatting artifacts for a new AO's preferred template is an hour not spent on the substantive security challenges that AI-enabled systems pose. Administrative burden is not a neutral cost. It actively displaces the technical attention that genuine security requires.

Automation efficiency is one of the core tenants of CSRMC. Some IC components have already demonstrated what this looks like in practice: for example, using ServiceNow to provide real-time monitoring and automated documentation generation, driven by the same configuration management tools the systems already require for operations. That approach proves the concept, but it also illustrates the risk of fragmentation. When individual components build bespoke automation stacks, the documentation they produce remains locked in unique formats that may not be transferred or reused across organizational boundaries. Rather than mandating uniformity in the underlying tools—which would suppress legitimate competition and innovation among DevSecOps pipeline vendors—OSCAL establishes a standardized output format that accommodates variation in the tools themselves while ensuring the compliance artifacts they produce are interoperable. Organizations can compete on the quality of their security scanning, configuration management, and continuous monitoring capabilities; what they cannot do is produce compliance documentation that only their own reviewers can read. The result is an ATO package that is automated and sharable while reducing duplication across the entire federal government. According to one practitioner, AI is now automating relevant information in knowledge bases full of project data to depict how each control is answered.

Vignette: Recommendation 2

Because all AOs now require OSCAL-formatted compliance submissions, your contractor no longer produces static artifacts for a single assessment event. Instead, they leverage their existing configuration management tools, DevSecOps pipelines, and AI agents to generate compliance evidence continuously, flowing directly into authorization repositories in a structured, machine-readable format. AOs, CIOs, and senior leadership gain real-time visibility into the critical controls through their own AI-enabled review tools, without waiting for a contractor to repackage documentation on demand. Continuous monitoring becomes substantively meaningful rather than nominally required: Stakeholders have live insight into software versions, software bills of materials, and control status as systems evolve.

Critically, the contractor built this capability before engaging with the government because OSCAL compliance was a known federal requirement from the outset, not a retroactive demand discovered mid-procurement. That up-front investment serves not just DoW AOs but also IC customers and every other federal agency, since it was mandated by OMB. The result is an ATO package that is automated and sharable while reducing duplication across the federal government.

Recommendation 3: Budget and Deploy Continuous, Automated AI Red Teams

Once systems meet critical security controls, program offices and AOs should transition from treating the remaining NIST controls as a compliance checklist to implementing continuous, automated AI red teaming through the system life cycle. Leveraging agentic AI for scalable, unannounced threat assessments provides real-time security insights while optimizing the deployment of scarce, highly skilled human cyber talent.

High-quality human red teams are difficult to scale because they require deep mission context, specific system knowledge, and up-to-date threat intelligence. Conversely, agentic AI teammates can execute realistic, on-demand threat assessments at scale. The federal government and DoW must harness this capability for internal continuous evaluation throughout development, testing, deployment, and operations. Conducting these continuous assessments and unannounced yields a much more accurate picture of a system's cyber resilience than previous methodologies.

Skilled offensive cyber operators and vulnerability researchers are rarely ATO process specialists. Currently, the scarcity of personnel possessing deep technical expertise, mission context, and RMF mastery creates cybersecurity compliance theater and delays. A strategic solution is to reskill the existing compliance workforce. Transitioning to AI-led red teaming does introduce risks, such as automation bias, cognitive off-loading, and AI accuracy issues. However, similar vulnerabilities already exist in the current system, where assessors lacking deep technical expertise often rely on checklists without understanding the technology or mission context. This risk can be actively managed by shifting the compliance community's focus to specialize in AI-generated security assessments and mitigating the limitations of various automated tools.

Because AI agents require a financial investment, programs must incorporate these costs into their overall budget plans, and Congress must appropriate the necessary funding. Several budgeting strategies can support this recommendation, including implementing a dedicated percentage tax on overall budget dollars specifically set aside for AI red teaming or requiring program managers to explicitly break out and include estimates for AI red teaming within their formal program estimates.

Vignette: Recommendation 3

Having already met the critical controls required by the DoW to obtain the ATO, your team is free to focus on the remaining threats most relevant to the mission. By leveraging your automated AI red teaming tools, you can rapidly prioritize the cybersecurity backlog based on actual threat criticality rather than context-free, legacy scans. Because Congress provided sufficient funding, this continuous evaluation begins early in the development environment. As a result, your assessors, AO, and CIO gain more frequent, meaningful, and actionable insights into the mission's true cybersecurity risk posture.

Recommendation 4: Experiment with AI Reciprocity Agents

A bottleneck to scaling technology across the DoW is the manual, inefficient process of security reciprocity. When a new unit adopts an existing system, they deploy it within their unique operational context. This variance can render the initial risk assessment of critical controls either highly relevant or completely outdated.

Furthermore, as programs leverage P-ATOs for cloud services—which is heavily emphasized in the FY2026 NDAA—acquiring and integrating those existing assessments is manual and complex. Current reciprocity requires security teams to manually request access, navigate disparate databases—such as GSA's FedRAMP Marketplace, DoW's eMASS, the IC's Xacta—and manually import data into the new ATO package. An AI reciprocity agent would be able to continuously monitor and pull new assessment reports to feed the AI red teams, which helps shift from static, point-in-time assessments to more timely threat information. Programs and AOs do not usually prioritize keeping up with every single new assessment report for every single cloud service or technology. They are typically dependent on other organizations to help notify them of those issues.

The AI reciprocity agent would continuously monitor underlying cloud services, and if a new vulnerability is discovered in the underlying P-ATO cloud service, the AI reciprocity agent would quickly feed the data to the system's local red teaming agent, allowing the local system to determine if the vulnerability actively threatens its specific architecture or if existing security measures render it ineffective.

While an AI reciprocity agent would dramatically accelerate potential vulnerability notifications, it would also introduce critical security risks that must be managed. The most pressing concern is the aggregation of data threat. Because this agent requires access to extensive vulnerability data across the federal government, it would become a high-value target where uncontrolled data aggregation could inadvertently expose systems' weaknesses. Additionally, the risks of AI red teaming are still a concern, such as automation bias, cognitive off-loading, and AI accuracy issues.

To manage these risks, a joint task force consisting of DoW, IC, and federal AI experts must govern and audit this capability. This body would enforce struct guardrails, which could include obfuscating specific system and unit names with shared datasets, implementing manual preview protocols for exporting outside of the system, and regularly red teaming the reciprocity AI itself to ensure it cannot be manipulated into leaking sensitive vulnerability data or turned against internal systems.

This recommendation should not be implemented until it has been attempted with low-impact, unclassified systems to validate the agent's accuracy and security guardrails. Some systems, such as Nuclear Command, Control, and Communications (NC3), should be permanently excluded from AI reciprocity agents. If the benefits outweigh the risks in the first limited phase, then expanding capabilities to higher classification levels should be implemented.

Vignette: Recommendation 4

Your drone swarm is deploying a system that leverages cloud services previously authorized by an IC AO. Instead of manually chasing down compliance packages and obtaining a static copy, your AI red team is connected to the AI reciprocity agent that has access to all the federal cyber compliance repositories. The agent securely accesses the IC's repository, pulls the relevant assessments, and automatically feeds your local AI red team. When a new vulnerability is discovered in the underlying cloud service, you do not have to wait for the next review, when you must manually pull any assessment data that may have changed. Your assessors, AO, and both the DoW and IC CIOs gain more frequent, meaningful, and actionable insights into the mission's true cybersecurity risk posture.

Recommendation 5: Increase Reciprocity and Cybersecurity Incentives

To accelerate technology scaling across the military and IC, the DoW and IC CIOs could mandate automatic reciprocity for defined system classes that meet standardized critical controls and possess a valid ATO at the appropriate classification level. Rather than requiring redundant, duplicative ATOs across the services and agencies, the default posture could shift to a "presumption of adequacy," backed by continuous automated AI red teaming and an explicit "no list" for specific types of platforms or missions.

Under the current policy, reciprocity guidance encourages AOs to reuse prior security assessments yet still mandates "careful consideration" for each system.[45] This soft requirement can result in a highly cautious, duplicative approval process where individual AOs generate redundant documentation and demand brand-new ATO packages for similar systems across different services and commands. The proposed shift replaces this compliance-heavy approach with automatic reciprocity for defined classes of software and AI systems that meet standardized critical controls. Under this new model, a valid ATO issued by any DoW or IC AO grants immediate, service-wide fielding rights. Rather than front-loading risk management through static paperwork, security is maintained dynamically after implementation via continuous AI red teaming. A strict "no list" would explicitly define the few high-risk missions or legacy technologies exempt from this automatic reciprocity.

The Defense Innovation Board (DIB), now part of the Science, Technology, and Innovation Board, previously recommended a more aggressive alternative: Skip pre-

authorization compliance requirements entirely, allow vendors to compete on performance and innovation, and impose financial liability on the back end if something goes wrong.[46] The appeal is real: lower barriers to entry, more competition, and faster access to commercial innovation. But for the military and IC CIOs, that risk is likely too great. A system that was never assessed prior to fielding and subsequently compromises sensitive data or mission-critical infrastructure cannot be rectified through a financial penalty. The ex post model may have merit in lower-risk commercial contexts, but it is unlikely to be adopted by organizations where the consequences for failure extend beyond financial loss to operational compromise and loss of life.

That said, the DIB's recommendation shares the same core concern as this one: that compliance burdens are too high, too slow, and too disconnected from actual security outcomes at scale. Instead of eliminating the initial security bar, the DoW and IC should set clear expectations for the initial ATO by publishing and standardizing critical controls and then adopt an ex post risk management model as the technology scales. For software-defined, edge-updatable systems, the risk of addressing vulnerabilities after fielding can be lower than the operational risk of delaying critical capabilities.

For the "no list," AOs should be required to actively lobby to get a system placed on list. This way, the DoW and IC can establish an organizational bias toward rapid scaling. The list could include systems like NC3 satellites, piloted next-generation fighter aircraft, nuclear submarines, and legacy systems utilizing unmodernized architectures.

This proposed reciprocity model is highly dependent on recommendation 3. AI red teaming would ensure that AOs and CIOs maintain visibility into active risks as the system scales. This is also the method to actually impose fines or hold personnel accountable. If programs become aware of heightened cybersecurity risks to the system and do not plan to address the issue, then the actions described by the DIB could be taken.

Finally, the military and the broader federal government need a more nuanced approach to cybersecurity incentives than the binary ATO status quo. Contracting mechanisms such as award feeds could incentivize contractors to respond rapidly to relevant threats discovered by the AI red teaming. AOs and CIOs could establish guardrails that automatically remove systems that fail to prioritize and resolve known vulnerabilities. Contractors and program managers could receive cybersecurity report

cards that follow their professional records, creating accountability that persists beyond any single program.

Vignette: Recommendation 5

After the drone swarm is fielded with the Air Force unit, the Army and Navy are cleared to field it for their missions. The Army and Navy AOs, along with the DoW and IC CIOs, have insight into their cybersecurity risks through the local AI red teams on each system. Rather than waiting for multiyear reassessment cycles, you leverage the timelier threat information to prioritize security vulnerabilities that affect the missions. If the risk increases for one of the systems and you do not take action, the AO or the DoW CIO can address the increased risk with the unit.

Who Should Establish the Policy?

The DoW CIO could direct recommendations 1 through 3 unilaterally within the military. The deputy secretary of war could enforce them at the senior leadership level. The IC CIO and the principal deputy director of national intelligence could pursue parallel action within IC authorities, but voluntary coordination between the DoW and IC has historically moved slowly on process questions that implicate institutional prerogatives. The highest impact path runs through OMB, whose authority over government information security could mandate transparency across the entire federal enterprise, including the IC, producing benefits at scale that department-level directives cannot achieve. Congress is the alternative forcing function: The same legislative mechanism that produced the FY2026 NDAA's accountability provisions could require publication of critical control lists and AO process requirements as a condition of authorization funding. If the DoW acts alone and the IC does not follow, the dual-authorization problem that affects military AI systems levering IC SCI data will remain.

Recommendation 4 requires action at a level above any individual department or agency. Solving the federal reciprocity problem at scale requires a common repository architecture, cross-organizational data governance, and a willingness to treat authorization documentation as shared infrastructure. No single CIO can compel that. OMB would need to establish the data governance framework and address the connectivity questions that a shared authorization repository raises. Congress may need to provide the directive authority and, where necessary, the funding to build and sustain the infrastructure. This is not a recommendation that can be implemented

incrementally by willing components; it requires coordinated top-down direction. Congress and OMB are the right actors to establish it.

Recommendation 5 applies to systems within the DoW and IC only. These systems carry out similar or complementary missions that are not found in other agencies across government. The DoW CIO and IC CIO would establish and maintain a joint "no list" of technologies and missions that will not receive automatic reciprocity.

OMB and Congress are the appropriate actors to mandate reciprocity across both organizations. Whether to establish a separate mandate to the broader civilian federal enterprise is another question. Federal agencies may have sufficiently different missions and risk environments that a uniform mandate is not appropriate.

Table 3 provides a summary of the recommendations, key actions, anticipated impacts, appropriate policymakers, and relative difficultly to implement.

Table 3. Summary of Recommendations

#	Recommendation	Key Actions	Impacts	Policymaker	Difficulty
1	**Publish critical controls and authorizing officials**	Mandate **public release** of AO-delegated authorities and critical controls.	AO matching is faster and increases the likelihood that the AO understands the technology. Schedule and cost risk is decreased for contractors and program managers since the compliance expectations are clear.	OMB, Congress, or DoW and IC CIOs	**Easy**
2	**Leverage AI to standardize to machine-readable formats**	Require **OSCAL**-formatted compliance outputs across all ATO submissions. Deploy **AI-enabled tools** to normalize documentation and auto-generate artifacts.	Reciprocity and automation increase. Competition of DevSecOps and compliance tools remains, but compliance is automated for auditors. Auditors are not required to be experts in every tool or technology and mission context.	OMB, Congress, or DoW and IC CIOs	**Medium**
3	**Budget for continuous, automated AI red teams**	Require budgeting for and deployment of **AI red teams** for continuous threat testing across the entire system life cycle.	AOs and assessors can implement more effective threat monitoring without having deep expertise in every threat or technology.	OMB, Congress, or DoW and IC CIOs	**Hard**

#	Recommendation	Key Actions	Impacts	Policymaker	Difficulty
			Program managers, contractors, and operational users can better prioritize the security backlog based on threats and mission context.		
4	**Experiment with AI reciprocity agents**	**Connect federal authorization repositories** to allow AI to retrieve appropriate assessments for reuse.	AI agents would have more up-to-date threat information, vast expertise on various types of technologies, and mission-specific context.	OMB or Congress	**Hard**
5	**Increase reciprocity and cybersecurity incentives**	Establish presumption of adequacy for systems with an ATO scaling across the DoW. Leverage other financial and performance incentives outside of the ATO process to help prioritize cybersecurity.	Technology scales rapidly across the military after initial ATO is achieved. Programs and contractors prioritize cybersecurity features due to the apparent fiscal impacts.	Congress or DoW and IC CIOs	**Easy**

Conclusion

Software now underpins nearly every aspect of modern warfighting, from intelligence analysis and command and control to autonomous systems and logistics. The federal cybersecurity authorization system, rooted in FISMA and implemented through the RMF, was designed for an earlier era of IT. As this paper has documented, that mismatch is not merely an administrative inconvenience. It is an operational risk. Civilians may have died due to critical software fixes stuck in the ATO process.[47]

Over the past decade, Congress and the DoW have taken meaningful steps to modernize the process through DevSecOps platforms, continuous authorization, FedRAMP standardization, and, most recently, the accountability mechanisms of the FY2026 NDAA. These reforms have produced real improvements, particularly for programs that have already made the up-front investment in DevSecOps infrastructure and have established AO relationships. The claims of dramatic authorization acceleration—ATOs in days, continuous authorization at the speed of release—reflect genuine achievements within that population. They do not reflect the experience of a program manager navigating the process for the first time with a novel architecture, raw intelligence data requirements, and no existing relationships with the officials whose approval they need.

This paper traced that experience from beginning to end. The friction points documented in the Background section reveal a process that has challenges at every step—from identifying an AO before step 0 begins to achieving meaningful reciprocity after an ATO is granted. The gap analysis that followed identified which points of friction reforms have focused on. While the reforms have good intentions, the resulting lack of observed change comes from the deeper process-oriented facets of RMF implementation across DoW and the inherent nature of navigating a disaggregated, large-scale, nonstandard enterprise such as the military.

The five recommendations that this paper proposes address each of these conditions in sequence. The authorization process was designed to ensure that systems are secure enough to operate, not to maximize the documentation produced in service of that determination. Transparency of process does not weaken security. Standardization does not eliminate judgment. Reciprocity does not remove oversight.

These recommendations carry real risk that must be managed: aggregation of vulnerability data, automation bias, cognitive off-loading, and AI accuracy limitations all require deliberate governance structures, phased implementation, and continuous human oversight. The current ATO system already accepts significant risk. Assessment

personnel rarely possess deep technical expertise in the latest innovative commercial software, understanding of the threats to that software, mission context, and RMF mastery simultaneously, which produces compliance theater rather than real security. AI red teaming and reciprocity agents offer an opportunity to more actively manage risk. By transitioning the compliance workforce from attempting to master every technology and threat to governing and auditing AI red teams and agents, the government can scale scarce expertise in ways human staffing models cannot. Some personnel will still need deep technical knowledge to train models, validate outputs, and build guardrails so that the red teams and agents do not turn against the mission systems. But the alternative of continuing to accept the compounding risks of an under-resourced, checklist-driven compliance system is not a winning risk management strategy.

The U.S. military's ability to compete depends not only on building better software than the adversary but on rapidly fielding it while still maintaining awareness of the latest threats and potential impacts to mission execution. The authorization process must support that mission. As currently designed, it too often does not. These recommendations offer a path toward a process that does, at a moment when both policy appetite and AI technology maturity make reform achievable.

Author

Major Katherine “Katie” L. Carroll is a former Space Force fellow at CSET.

Acknowledgments

This project would not have been possible without the insights of experienced and passionate ATO experts who generously shared their time and expertise. I am especially grateful to Brian Panello, Michael Medgyessy, and Matthew Connor, who gave their time across multiple meetings and email exchanges. Your willingness to engage so openly allowed me to connect ideas and perspectives in ways I could not have managed alone. Any errors, mischaracterizations, or oversimplifications that remain are entirely my own.

I am also grateful to the following individuals who shared their experiences and perspectives: Sarah Pearson, Brandi Pickett, Jason Ingalls, Bryon Kroger, Alexa Sperling, Remington Brown, James Davidson, Jonathan Ng, and David Raley.

I would not have seen this project through without the support of my CSET research advisor, Emelia Probasco. Her confidence in my work helped bring clarity to a complex topic and helped me flesh out recommendations that could make an impact at scale. Drew Lohn’s willingness to challenge my thinking and identify where my recommendations fell short made the final product considerably stronger. The intellectual culture you both cultivate at CSET is genuinely distinctive, and the leadership you model is something I aspire to.

Thanks also to Igor Mikolic-Torreira, Lauren Kahn, John Bansemer, Peter Mastro, and Jeffrey Demetrelis for their feedback, editorial review, and support.

Disclaimer

The views expressed in this report are those of the author and do not necessarily reflect the official policy or position of the U.S. Department of the Air Force, the U.S. Department of War, or the U.S. government.

Document Identifier: doi: 10.51593/20250028

Endnotes

[1] "CITI Hearing: Too Critical to Fail: Getting Software Right in an Age of Rapid Innovation," House Armed Services Committee, 118th Congress, March 13, 2024, https://armedservices.house.gov/calendar/eventsingle.aspx?EventID=3558.

[2] Bonnie Evangelista, host, *Defense Mavericks*, podcast, "The Power of Continuous Software Delivery in Defense with Bryon Kroger," accessed April 23, 2024, www.defensemavericks.com/the-power-of-continuous-software-delivery-in-defense-with-bryon-kroger/.

[3] Defense Innovation Board, *Scaling Nontraditional Defense Innovation* (Washington, D.C.: DoD, January 2025, https://stib.cto.mil/wp-content/uploads/2026/01/2025-2_DIB-ScalingNontraditionalDefenseInnovation_250113PUBLISHED_9ee4ae.pdf.

[4] House Armed Services Committee, "Too Critical to Fail."

[5] Joe Tidy, "Iran War: What Role Is Cyber Warfare Played in Iran?," BBC, March 12, 2026, www.bbc.com/news/articles/c5yr0576ygvo.

[6] Brian Humphreys, "Attacks on Ukraine's Electric Grid: Insights for U.S. Infrastructure Security and Resilience" (Congressional Research Service, May 17, 2024), www.congress.gov/crs-product/R48067.

[7] Sam Bresnick, Ngor Luong, and Kathleen Curlee, "Which Ties Will Bind?" (CSET, February 2024), https://cset.georgetown.edu/publication/which-ties-will-bind/.

[8] Evangelista, "The Power of Continuous Software Delivery in Defense."

[9] Jason Miller, "OMB Sets 2017 as Deadline to Move to Dynamic Cybersecurity," Federal News Network, November 19, 2013, https://federalnewsnetwork.com/technology-main/2013/11/omb-sets-2017-as-deadline-to-move-to-dynamic-cybersecurity/.

[10] DoW, "Remarks by Secretary of War Pete Hegseth at SpaceX," January 12, 2026, www.war.gov/News/Transcripts/Transcript/Article/4377190/remarks-by-secretary-of-war-pete-hegseth-at-spacex/.

[11] Jamie Porchia and Daniel J. Finkenstadt, "Unpacking the Authority to Operate Process: Implications for the DoW," *Proceedings of the Twenty-Second Annual Acquisition Research Symposium and Innovation Summit* 2 (May 8, 2025),

https://dair.nps.edu/bitstream/123456789/5375/1/SYM-AM-25-342.pdf; Marc B. Bucks and Daphne Williams, "Maintaining a Competitive Advantage: An Analysis of the Efficiency and Effectiveness of the Marine Corps' Assessment and Authorization Process" (thesis, Naval Postgraduate School, September 2021), https://apps.dtic.mil/sti/trecms/pdf/AD1164217.pdf.

[12] GAO, *High-Risk Series: Urgent Action Needed to Address Critical Cybersecurity Challenges Facing the Nation* (Washington, D.C.: GAO, June 13, 2024), www.gao.gov/products/gao-24-107231; GAO, *Cybersecurity: OMB Should Update Inspector General Reporting Guidance to Increase Rating Consistency and Precision* (Washington, D.C.: GAO, March 31, 2022), www.gao.gov/products/gao-22-104364; GAO, *Cloud Security: Federal Authorization Program Usage Increasing, but Challenges Need to Be Fully Addressed* (Washington, D.C.: GAO, January 18, 2024), www.gao.gov/products/gao-24-106591; GAO, *Cloud Computing Security: Agencies Increased Their Use of the Federal Authorization Program, but Improved Oversight and Implementation Are Needed* (Washington, D.C.: GAO, December 12, 2019), www.gao.gov/products/gao-20-126.

[13] Federal Information Security Modernization Act of 2014, Pub. L. No. 113-283, 128 Stat. 3073 (2014), www.congress.gov/bill/113th-congress/senate-bill/2521.

[14] NIST, "2.5 Authorization Boundaries," in *Risk Management Framework for Information Systems and Organizations*, NIST SP 800-37, rev. 2 (NIST, December 2018) https://nist-sp-800-37-r2.bsafes.com/docs/chapter-two/2-5-authorization-boundaries/.

[15] Dominick A. Fiorentino and Meghan M. Stuessy, *Chief Information Officers (CIOs): Agency Roles and Responsibilities* (Washington, D.C.: Congressional Research Service, August 1, 2024), www.congress.gov/crs-product/R48147.

[16] FedRAMP, "FedRAMP Built a Modern Foundation in FY25 to Deliver Massive Improvements in Fy26," news release, September 30, 2025, www.fedramp.gov/2025-09-30-fedramp-built-a-modern-foundation-in-fy25-to-deliver-massive-improvements-in-fy26/.

[17] GAO, *Cloud Security*.

[18] FedRAMP, "FedRAMP Built a Modern Foundation."

[19] Department of Defense, *Department of Defense Chief Information Officer Desk Reference*, vol. 1, *Foundation Documents* (Washington, D.C.: DoD, August 2006), https://dodcio.defense.gov/portals/0/documents/ciodesrefvolone.pdf.

[20] DoD, *Cybersecurity Reciprocity Playbook* (Washington, D.C.: DoD, March 2024), https://dodcio.defense.gov/Portals/0/Documents/Library/(U) 2024-01-02 DoD Cybersecurity Reciprocity Playbook.pdf.

[21] DoD CIO, *Risk Management Framework for DoD Systems*, DODI 8510.01, July 19, 2022, www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/851001p.pdf.

[22] Under Secretary of Defense (Acquisition and Sustainment), *Operation of the Software Acquisition Pathway*, DODI 5000.87, October 2, 2020, www.esd.whs.mil/Portals/54/Documents/DD/issuances/DoDi/500087p.PDF.

[23] Office of the Chief Information Officer, "Air Force (AF) Risk Management Framework (RMF) Information Technology (IT) Categorization and Selection Checklist (ITCSC)," Department of the Air Force, December 3, 2019, www.afacpo.com/AQDocs/AF_Information_Technology_Categorization_and_Selection_Checklist_v2.0.pdf.

[24] GAO, *Weapon Systems Cybersecurity: Guidance Would Help DOD Programs Better Communicate Requirements to Contractors* (Washington, D.C.: GAO, March 4, 2021), www.gao.gov/products/gao-21-179.

[25] Fiorentino and Stuessy, "Chief Information Officers (CIOs)."

[26] NIST, *Security and Privacy Controls for Information Systems and Organizations*, NIST Special Publication 800-53, rev. 5. (Washington, D.C.: Department of Commerce, September 2020), https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf.

[27] Hannah Hunt, "The Pentagon's Software Approval Process Is Broken. Here's How to Fix It," Atlantic Council, June 4, 2025, www.atlanticcouncil.org/blogs/new-atlanticist/the-pentagons-software-approval-process-is-broken-heres-how-to-fix-it/.

[28] GAO, *Cloud Security*.

[29] Mikayla Easley, "Inside the Pentagon CIO's Push to Overhaul Antiquated Software Acquisition Practices," *DefenseScoop*, June 9, 2025, https://defensescoop.com/2025/06/09/katie-arrington-swft-software-fast-track/.

[30] Jason Miller, "Arrington Kicks Off Effort to Eliminate RMF for DoD Software," Federal News Network, May 5, 2025, https://federalnewsnetwork.com/defense-main/2025/05/arrington-kicks-off-effort-to-eliminate-rmf-for-dod-software/.

[31] DoD, *Cybersecurity Reciprocity Playbook.*

[32] Nick Sinai, "ATO in a Day," Medium, July 25, 2017, https://nicksinai.medium.com/ato-in-a-day-fd53f0ef4a5d; Second Front, "Obtain an ATO in as Little as 90 Days," accessed September 16, 2025, www.secondfront.com/solutions/DoD-accreditations/; Ross Gianfortune, "Marine Corps Operation StormBreaker Slashes Software Delivery Timelines by 17x," GovCIO Media & Research, July 1, 2025, https://govciomedia.com/marine-corps-operation-stormbreaker-slashes-software-delivery-timelines-by-17x/.

[33] William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, 134 Stat. 3388 (2021), www.congress.gov/bill/116th-congress/house-bill/6395.

[34] GAO, *Weapon Systems Cybersecurity*.

[35] James M. Inhofe National Defense Authorization Act for Fiscal Year 2023, Pub. L. No. 117-263, 136 Stat. 2395 (2022), www.congress.gov/bill/117th-congress/house-bill/7776/text.

[36] Servicemember Quality of Life Improvement and National Defense Authorization Act for Fiscal Year 2025, Pub. L. No. 118-159, 113 Stat. 1773 (2024), www.congress.gov/bill/118th-congress/house-bill/5009/text.

[37] National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119-60, 139 Stat. 718 (2025), www.congress.gov/bill/119th-congress/senate-bill/1071/text/eah.

[38] FedRAMP, "FedRAMP Built a Modern Foundation."

[39] "What Is a Continuous ATO?," Rise8, accessed April 19, 2026, https://rise8.us/resources/what-is-a-continuous-ato.

[40] SAF/CN, March 18, 2019, Fast-Track Authorization to Operate (ATO), https://federalnewsnetwork.com/wp-content/uploads/2019/04/AF-fast-track-ATO-memo-march-2019.pdf.

[41] DoD, *Continuous Authorization to Operate (cATO) Evaluation Criteria* (Washington, DC: DoD, May 29, 2024), https://dodcio.defense.gov/Portals/0/Documents/Library/cATO-EvaluationCriteria.pdf?ver=A8tLlfPjmp3RpemU6JOhJw==.

[42] DoW, "Department of War Announces New Cybersecurity Risk Management Construct," news release, September 24, 2025, www.war.gov/News/Releases/Release/Article/4314411/department-of-war-announces-new-cybersecurity-risk-management-construct/.

[43] Cole McFaul, Sam Bresnick, and Daniel Chou, "Pulling Back the Curtain on China's Military-Civil Fusion: How the PLA Mobilizes Civilian AI for Strategic Advantage" (CSET, September 2025), https://cset.georgetown.edu/publication/pulling-back-the-curtain-on-chinas-military-civil-fusion/.

[44] NIST, "OSCAL: The Open Security Controls Assessment Language," Department of Commerce, accessed April 19, 2026, https://pages.nist.gov/OSCAL/.

[45] DoD, *Cybersecurity Reciprocity Playbook.*

[46] Defense Innovation Board, *Scaling Nontraditional Defense Innovation*.

[47] Evangelista, "The Power of Continuous Software Delivery in Defense."