Issue Brief

Putting Teeth into AI Risk Management

Lessons from Cybersecurity Procurement Rules and Practices

Author
Matthew Schoemaker

CSET CENTER for SECURITY and EMERGING TECHNOLOGY

May 2024

Executive Summary

President Joe Biden's recent signing of a sweeping executive order aimed at increasing governance of artificial intelligence in the federal government brings urgency to the creation and implementation of AI risk management standards and federal procurement guidelines.[1] The Office of Management and Budget (OMB) quickly followed with its guidance to departments and agencies, which includes AI minimum risk standards and their incorporation into federal contracts.[2] A looming challenge is how the government can best utilize federal procurement rules, requirements, and practices to ensure supplier compliance with AI development best practices.

The federal government often utilizes its significant purchasing power to incentivize and enforce policies among its industrial base, making compliance a condition of being awarded government contracts. The U.S. government's position as a major customer of many top companies has effectively made its cybersecurity framework the "de facto standard" that has been adopted by governments and industries worldwide.[3] The effectiveness of procurement rules in increasing the use and adoption of best practices in the case of cybersecurity has led members of Congress and industry leaders to cite it as an example to follow for AI risk management enforcement.

The current evolution of AI risk management frameworks, the corresponding legislation driving their development and use in the government, and the calls for their inclusion in federal procurement regulations are similar to the conditions that drove the creation of cybersecurity frameworks and federal procurement rules. For this reason, the federal government's adoption and implementation of procurement rules to enforce cybersecurity standards within its supplier base provides a blueprint for AI and can help forecast upcoming difficulties. These previous lessons and challenges arising from the implementation of cybersecurity procurement rules include:

1. Difficulty balancing the level of risk management to the level of risk impact.

2. Difficulty balancing trust and verification in assessment requirements.

3. Difficulty in oversight and enforcement of workforce preparation and training.

4. Concerns about third-party auditing and government oversight.

5. Use of procurement rules to enforce incident reporting and sharing.

Using these lessons as a guide, this paper provides the following recommendations for policymakers looking to institute AI procurement practices and standards:

1. Develop standards to assess the level of risk and potential impacts of AI systems. Establish categories to differentiate the levels of risk AI systems pose and develop the appropriate risk management practices required for each category.

2. Base the level of requirements verification on the overall risk of the system. Compliance audits are costly, and therefore the federal government should utilize risk categories to determine which systems require compliance auditing.

3. Mandate and provide training on AI risk management standards for the federal acquisition workforce.

4. Leverage third-party auditors to support assessments of supplier compliance with AI risk management standards. This would solve labor limitations and skills gaps in the federal workforce, but it is important that final approval decisions rest with the government. Establish an AI standards center of excellence to provide government oversight and support compliance assessments.

5. Use contracting rules to incentivize and, when necessary, compel government suppliers to comply with AI incident reporting and cross-agency sharing.

The recommendations provide implementation guidance on how to avoid missteps of the past while also enabling timely adoption of best practices. Oversight and enforcement of supplier compliance with AI risk management standards will require a significant effort on behalf of the government, one that should be informed by the historical experiences in cybersecurity and that is tailored to meet the specific demands of AI technologies. These recommendations can help guide the establishment of effective procurement rules, practices, and enforcement infrastructure to best ensure AI risk management compliance and mitigate the realization of AI harms.

Table of Contents

Introduction

President Joe Biden's October 2023 signing of a sweeping executive order aimed at increasing governance of artificial intelligence in the federal government introduced urgency to the creation and implementation of AI risk management standards and federal procurement guidelines.[4] The Office of Management and Budget (OMB) quickly followed with its guidance to departments and agencies mandating the inclusion of AI risk standards into future federal contracts.[5] As agencies look to implement the executive order and OMB guidance, a looming question is how the federal government can best utilize federal procurement rules, requirements, and practices to put teeth into AI risk management standards and ensure supplier compliance.

Congress, industry, and academia have made similar calls for AI regulation and procurement rules and have cited earlier cybersecurity guidelines as the example to follow.[6] The evolution of federal cybersecurity procurement regulations provides a relevant case study to identify and forecast the challenges that AI compliance mandates may face. This document discusses what AI risk management regulations exist or have been proposed, why cybersecurity is a useful precedent for AI regulation, lessons learned from cybersecurity mandates in federal procurement, and corresponding recommendations for how to best utilize federal procurement to mandate AI risk management standards compliance.

Current and Proposed AI Regulations and Rules

The U.S. government's methods of enforcing AI risk management practices to date have been only voluntary. The Biden administration secured voluntary commitments from some of the top AI companies to help advance the development of safe, secure, and trustworthy AI, but the commitments lack enforcement mechanisms for the government to hold these companies accountable.[7] The more recent executive order and OMB guidance signal an intent to place greater controls on the management of AI risks to ensure that the government's suppliers are mitigating harms and using AI responsibly.

The path to AI risk management requirements, for the federal government and its suppliers, flows from legislation and executive orders. The 2020 National Artificial Intelligence Initiative Act directed the National Institute of Standards and Technology (NIST) to create an AI risk management framework (RMF) and also directed OMB to provide a plan for how federal agencies will responsibly develop and acquire AI technologies.[8] The 2022 AI Training Act directed the establishment of a training program for acquisition professionals on how to effectively procure and manage AI systems.[9] More recently, the Biden administration released its executive order on AI to ensure safe, secure, and trustworthy development and use of AI systems by the federal government.

The latest executive order directed OMB to develop risk management practices and states that the office "shall develop an initial means to ensure that agency contracts for the acquisition of AI systems and services align with the guidance."[10] OMB followed this direction by publishing guidance that sets baselines for AI risk management practices and requires those practices in federal contracts.[11] Further standards could come from the passing of legislation that would mandate the use of NIST's AI RMF across federal agencies and suppliers.[12] OMB's release of minimum risk management practices and the more thorough NIST RMF are intended to help mitigate harms associated with AI development and use. The creation of oversight and enforcement mechanisms, likely through contract requirements, will be critical for these policies and standards to achieve their risk management goal. Fortunately, this is not the first time the federal government has sought to enforce technical standards through procurement regulations.

Cybersecurity Risk Management in Contracts as a Case Study for AI

The emergence of the internet and information technologies in the 1990s and 2000s forced the government to find ways to adopt these technologies and led to cybersecurity practices to mitigate new risks. A similar challenge is facing the federal government as AI technologies are emerging and providing new opportunities and risks. Akin to the guidance of OMB to require AI risk management practices in federal contracts, officials utilized procurement rules to direct federal suppliers' compliance with cybersecurity practices. The challenge of mitigating risks from an emerging technology and the approach of enforcing suppliers' compliance make cybersecurity a compelling case study for forecasting challenges in oversight and enforcement of AI standards.

As an enforcement mechanism, procurement guidelines are well suited to mandate compliance with federal standards by government suppliers. They can also incentivize and influence behaviors across the private sector. The U.S. government awards over $760 billion in government contracts to suppliers each year, with $3.3 billion going to awards for AI capabilities in 2022.[13] To access this funding, federal suppliers must adhere to the government's requirements. These suppliers include traditional defense contractors (e.g., Boeing, RTX, Lockheed Martin) and commercial companies such as Microsoft, Google, and 3M. Meeting federal requirements can have positive downstream benefits because, as companies develop business processes and practices to meet government procurement requirements, they often pass these requirements onto the commercial side of their business as well. The U.S. government has used its position as a significant customer to influence the behaviors of private companies, including in cyber, and industry leaders have cited that the government could do the same in AI risk management.[14]

In 2002, Congress passed the Federal Information Security Management Act (FISMA), which:

1. Directed federal agencies to implement information security programs that supported operations and assets of the agency, including those managed by contractors or other sources.

2. Mandated federal reporting and evaluations of the security implementations.

3. Assigned NIST the responsibility to develop the standards, guidelines, and methods for securing information systems.

The executive branch issued its own policies that established standards and controls to manage cyber risks within the federal supply chain. Notably, the government enforced these controls through procurement guidelines that made it mandatory for suppliers to institute specific information security controls in accordance with NIST standards and to comply with incident-reporting requirements.

The U.S. government's position as a major customer of many top companies has effectively installed the federal cybersecurity framework, developed by NIST, as the "de facto standard" adopted worldwide.[15] As members of Congress and industry leaders have argued, the effectiveness of procurement rules in increasing the use and adoption of cybersecurity frameworks makes this example a strong one for AI risk management enforcement to emulate.[16]

Cybersecurity Procurement Challenges and Lessons for AI

Although cybersecurity standards and procurement rules have been noted as a precedent for AI risk management, enforcement of these rules has not come without issues. These implementation challenges include:

1. Difficulty balancing the level of risk management with the level of risk impact.
2. Difficulty balancing trust in vendors' commitment to risk management practices with the need for government verification.
3. Difficulty in oversight and enforcement of workforce preparation and training.
4. Concerns about third-party auditing and government oversight.
5. Use of procurement rules to incentivize incident reporting and sharing.

Acknowledging these challenges may help regulators and standard setters forecast challenges in AI risk management. This section provides background information on each of these implementation challenges and a corresponding recommendation for how best to overcome them in the case of AI.

Balancing the Level of Risk Management with the Level of Risk Impact

Ensuring that the government did not create unnecessarily onerous requirements for every government and vendor information system was an early best practice adopted for managing cyber risks. The 2002 FISMA legislation directed NIST to develop guidelines and standards for securing federal information and information systems, including:

1. Standards to define categories for all information systems and information according to a range of risk levels.
2. Guidelines recommending the types of information and information systems to be included in each defined risk category.
3. Minimum information security requirements for information and information systems in each such category.

NIST, in adherence with this directive, developed standards for the federal government to follow (Table 1) to determine the necessary level of risk mitigation and to enable security controls to be tailored to different risk levels.[17]

Table 1: Summary of the Documents NIST Developed in Compliance with FISMA Directives

NIST Standards Document	Purpose
Federal Information Processing Standard (FIPS) Publication 199	Categorizes information and information systems by impact levels (low, moderate, or high) based on the potential impact to an organization should there be a breach of confidentiality, integrity, or availability.
FIPS Publication 200	Specifies minimum security requirements for federal information and information systems across 17 security-related areas to achieve adequate security according to the risk levels identified by FIPS 199.
Special Publication 800-53	Defines the security and privacy controls for federal information and information systems to manage risks and ensure the confidentiality, integrity, and availability of information.

Managing all risks with the same care is neither efficient nor effective. Requiring the most stringent standards in all scenarios is an inefficient use of resources, and using only the minimum standards in all scenarios is insufficient. Security categories and requisite controls were created so that the risk management prescription is tailored to match the level of potential impact of the risk being realized.

Recommendation: Develop standards to assess the level of risk and potential impacts of AI systems. Establish categories to differentiate the levels of risk AI systems pose and develop appropriate risk management practices for each category.

Not all compromises of information systems incur the same risk or require the same level of security. Similarly, not every AI system will bring about the same level of risk or require the same level of risk management controls. OMB has taken a strong first step

in the recently released minimum compliance practices for systems that could impact the safety and rights of the public.[18] Looking beyond these minimum practices, the more thorough NIST AI RMF standards that members of Congress have advocated for as a federal requirement are currently voluntary.[19] To expand the use of the RMF as a requirement for federal agencies and suppliers, NIST will need to revise it or create a new standards document with mandatory requirements and assessment.

Additionally, the federal government should develop a standard for categorizing systems so that controls are tailored based upon risk. OMB's guidance establishes two categories of AI systems requiring risk management: systems impacting safety and systems impacting human rights. Again, this is a good first step by OMB, but two categories are unlikely to provide adequate distinctions across the range of potential AI systems.

An example of where the OMB standards do not provide adequate risk differentiation is when AI is used to make decisions in the movement of physical systems. OMB's standards impose the same level of risk management to an AI-enabled sorting arm on an assembly line and an AI-enabled seeker on a missile. While the OMB's proposed minimum standards are appropriate for both cases, additional risk management measures are needed for higher risk systems.

OMB does not ignore the need for additional risk management practices for certain systems, and in the guidance it recommends that agencies look to the NIST AI RMF and other frameworks. This is where the government lacks teeth in mitigating risk, as the NIST RMF is a voluntary document and does not impose risk mitigation requirements or categorization of risk levels. One way OMB could move toward a differentiated risk level approach is following the European Union (EU) AI Act's regulatory framework. The EU AI Act establishes four risk levels (minimal, limited, high, and unacceptable), which then drive the risk management requirements.[20] Requiring the NIST RMF's mapping and measuring of AI systems' risks would also enable the categorization of risks and enable a prescription of associated mitigation and management practices.

Balancing Trust in Vendors' Commitment to Risk Management Practices with Government's Need for Verification

The Department of Defense (DoD) was one of the first federal agencies to include enforcement of cyber requirements in its procurement rules. However, these rules have undergone multiple revisions to balance the amount of trust placed in supplier self-certification versus the need for extensive security audits. Initial rules for cybersecurity compliance merely required vendors to self-attest to being cyber secure.[21] A 2019

audit of nine federal contractors found that they were not meeting many of the security controls, including some of the most basic ones.[22] This audit highlighted that self-certification was insufficient to ensure compliance with the required security practices.

Following the 2019 audit, the DoD strengthened procurement requirements with the creation of three assessment levels. These levels correspond to the amount of trust the government will have in the supplier's certified compliance with cybersecurity requirements following the assessment.[23] The three cybersecurity confidence levels are:

1. Low confidence: supplier self-assessment.

2. Medium confidence: DoD review of supplier system security plans and supplier interview.

3. High confidence: on-site validation of security plan implementation.[24]

The DoD also proposed a rule that would require on-site validation and pre-auditing of supplier infrastructure before any contract award. This program, the Cybersecurity Maturity Model Certification (CMMC), would have driven most suppliers to require "high confidence" certification and accelerated suppliers' timelines for compliance and certification before even bidding on a contract.[25] Federal suppliers and government agencies pushed back on the bureaucratic burden CMMC would entail and raised concerns about the requirements for third-party auditing.[26] Following the pushback, the government retreated to the original three assessment tiers and began work on a revised CMMC assessment strategy. CMMC's revised strategy plans reduced the number of vendors requiring auditing from 221,286 to the 78,085 designated supplier systems that would host special-interest or sensitive but unclassified information.[27]

The latest CMMC rule provides insights into the cost burden of the assessment program on federal suppliers and the government, but notably it does not cite the costs of implementing the security controls to pass the assessment. The annual costs to the public sector are estimated to be $4 billion, as well as a DoD cost of $10 million.[28] Additionally, the draft rule shows the difference in costs to suppliers of self-assessment versus auditing and certification. Self-assessments were determined to cost a company approximately $14,300 annually, compared to $37,000 for assessment and certification. Over 80,000 companies fall within this security construct, and the government is planning to require almost all to have a certification, incurring an additional $1.8 billion in annual assessment costs. The DoD has spent years trying to balance cybersecurity requirements and assessments in its supplier base with the cost

burden of implementing the program, but industry continues to disagree with the balance in the rules being proposed.[29]

Though the DoD limited the mandate for on-site assessments to critical programs, challenges in ensuring compliance continue. A 2022 DoD report on audits for 117 federal suppliers identified hundreds of security requirements that were still not satisfied.[30] That same year, in a survey of 300 defense suppliers, 70 percent claimed compliance via self-assessment, but only 13 percent met the minimum requirements of compliance.[31] The DoD has taken the lead in pushing its supplier base to be compliant with the federally mandated cybersecurity standards, but its need to balance supplier audits with supplier burdens, as well as with its own labor and resource limitations, hinders progress.

Finding the right balance of trust and verification of its suppliers remains a challenge for the DoD. While having every supplier audited is likely the best way to ensure compliance, a balance needs to be struck between potential risks and the costs of completing audits. DoD's experiences in enforcement through contract mandates and audits highlight the potential challenges with similar AI mandates.

Recommendation: Compliance audits are costly, and therefore the decision to audit should be based on the system risks. Utilize the developed risk categories to guide which category of systems would require compliance auditing.

The level of auditing and assessment of suppliers is a critical consideration when creating procurement rules that are enforceable and effective. The current OMB guidance states that a federal AI council will develop "a list of recommended documentation that should be required from a selected vendor in the fulfillment of a federal AI contract."[32] Just as there are varying levels of audit requirements for cyber compliance, the government should consider requiring an audit of a company's AI risk management standards implementation for certain levels of risk. While the documentation for OMB is being developed, the AI council should recommend the levels and types of risk that would require further auditing and verification.

Appropriately Preparing the Acquisition Workforce

The complexity of enforcing cybersecurity standards highlights the need for a competent and trained workforce to implement federal procurement requirements. Though procurement requirements for cybersecurity practices have been in place for several years, the federal government is still struggling with its own workforce's ability to manage suppliers. The DoD's rollout of cybersecurity procurement requirements,

specifically to safeguard controlled unclassified information (CUI), highlights the impacts of "late to need" training and the lack of training enforcement.

To effectively enforce and manage a supplier's requirement, the federal government must first ensure it has an adequately trained workforce. A 2019 DoD inspector general (IG) report on implementation of cybersecurity risk management requirements found that contractors were improperly implementing security standards and that the government was not consistent in its oversight of the program.[33] The findings included government documents not having appropriate security markings to inform contractors of handling requirements, inconsistent recording of information exchanges, and lack of verification of contractors' network security. In 2020, following this report, the DoD introduced organizational instruction and annual training requirements for federal employees to better manage these programs.[34] While the federal requirement for suppliers to comply with cybersecurity standards started in 2016, and was mandatory by 2017, DoD workforce instruction and training was not codified until 2020.

The training must not only be timely to support the implementation of procurement requirements, but it must also be effective and enforced. A 2023 DoD IG report on the effectiveness of a CUI cybersecurity training program found that the department developed training guidance but did not effectively oversee implementation, that DoD personnel did not consistently complete training, and that DoD personnel did not effectively oversee contractors' completion of their required training.[35] Within the DoD's vast workforce, the prevalence of cyber and information security in the majority of employees' daily work is significant, and responsible handling is imperative. The IG's findings highlighted the need for timely training material, enforcement of training, and oversight of the execution of the trained skills—all of which are necessary to ensure the government is doing its part in enabling and overseeing the cybersecurity practices of its suppliers.

The implementation of CUI cybersecurity requirements on suppliers required a significant level of effort by government agencies to train their employees to execute and enforce the program. The federal government and its suppliers are still struggling to work together to effectively implement these requirements, and a significant impediment is having a knowledgeable and ready government workforce to do its part. This highlights the need for timely and enforced training of the federal acquisition workforce to provide oversight of supplier cybersecurity compliance.

Recommendation: Mandate and provide training on AI risk management standards for the federal acquisition workforce.

In preparation for mandating AI risk management procurement requirements, it is important to learn from the lessons from CUI cybersecurity enforcement with suppliers. These include timely training for the workforce that is procuring AI systems and prioritizing and enforcing workforce training requirements. The prevalence and complexity of AI technology in future DoD procurements will require greater partnership between acquisition offices and assessment organizations and should require AI and risk management training for any office that is contracting the use of this technology. This need for training has already received legislative support in the AI Training Act, which directed OMB to develop and deliver training for the acquisition community and was a major point of emphasis in the latest AI executive order.[36] Implementation of a training program prior to standards enforcement will be critical in ensuring that the federal government is prepared to execute its oversight role.

Third-Party Auditing Concerns

The utilization of third-party auditing to perform independent compliance assessments is common in many industries, including health care, food, and cybersecurity. However, the federal government has faced challenges with third-party auditing of supplier compliance in managing sensitive but unclassified information. The government has seen much wider adoption of third-party auditing in assessments of federal cloud service providers' security requirements under the Federal Risk and Authorization Management Program (FedRAMP). The critical differences between the comprehensive cybersecurity audits and the cloud service provider audits were in the number of audits required and the amount of decision authority the government delegated to the third-party auditors. A similarity is that both efforts required the establishment of a government office to conduct and oversee supplier assessments.

The DoD's CMMC program, along with its attempt to mandate on-site pre-auditing of all prospective suppliers' cybersecurity compliance, tried to resolve concerns about the workforce needs of the program through the use of certified third-party auditing organizations (C3PAOs). While this was intended to follow industry best practices, this pivot to non-governmental organizational assessments received significant pushback. Of note, before taking the mantle as Secretary of the Air Force, Frank Kendall wrote an op-ed describing challenges with CMMC. His primary concern was that the C3PAO was taking over an inherently governmental function of deciding on a contractor's qualifications to bid on a contract.[37] This type of feedback, along with 850 public

comments, led to the elimination of 65 percent of cases where C3PAOs would be used in the audits and expanded government oversight of C3PAO practices.[38]

The federal government also has specific cybersecurity requirements for cloud services before their procurement and authorized use. These security standards were established in the form of the FedRAMP, which is a program that provides the process and standards by which a supplier's cloud offering can receive a government authority to operate (ATO) certification.[39] The ATO allows a cloud service to be used to store and process government information. As part of this certification, a third-party auditor assesses a potential cloud service provider and delivers a report to the company and the government. This report is a critical to the final determination of whether to provide an ATO, but the decision authority rests with the government.

Additionally, requiring every DoD office to have the skilled personnel needed to perform assessments of the highly technical cybersecurity controls levied by procurement requirements was not realistic and led to ineffective oversight and inefficient execution.[40] The DoD found that it had to assign responsible organizations to specialize in cybersecurity assessments. It assigned the Defense Contract Management Agency the responsibility for the assessment of CUI security, leading to the establishment of the Defense Industrial Base Cybersecurity Assessment Center, and tasked the FedRAMP office with overseeing cloud security assessments.[41] The establishment of organizations responsible for compliance oversight was necessary so that all DoD offices would have access to this expertise within the government.

The use of third-party auditors helps fill a need for qualified auditors when the government may not have sufficient expertise. The DoD has investigated the opportunity of using third-party auditors in cybersecurity audits but has decided to significantly limit their use and authority, while also ensuring maximum government oversight of their assessments. Should the federal government look to augment the workforce necessary to perform compliance assessments of AI risk management practices, it should ensure ATO or similar approval decisions are maintained as an inherently governmental function.

Recommendation: Leverage third-party auditors to support assessments of supplier compliance with AI risk management standards. This would solve labor limitations and skills gaps in the federal workforce, but it is important that final approval decisions rest with the government. Establish an AI standards center of excellence to provide government oversight and support compliance assessments.

AI, as an emerging technology, is a field in which the government does not have an extensive amount of workforce expertise. The workforce needed to accomplish the task of auditing contractual requirements is an important consideration in setting those requirements. It is not difficult to foresee a future where federal procurement policy mandates risk management standards for AI systems and finds a similar lack of compliance to what was witnessed in cybersecurity. In the cyber case, poor compliance led the federal government to seek greater oversight and enforcement through on-site auditing.

Rather than repeating past mistakes and trying to audit all uses of AI, the government can develop standards for classifying risk levels of AI systems and only utilize audits and third-party auditors in support of certain risk cases. In the case of cybersecurity, the use of third-party auditors was successful in auditing individual products from cloud service providers. AI systems are likely to be delivered as a product or software program offering, and because the assessment would be limited to the single AI system offering and not comprehensive company practices, third-party auditing could be a feasible and effective option. Finally, whatever use of third-party auditors is decided, the decision authority over whether a company has met its compliance requirements should remain a government decision and not delegated to a non-governmental entity.

Incident Sharing and Reporting Enforcement

Incident reporting is critical in managing and mitigating risks from cybersecurity vulnerabilities. While reporting of cyber incidents has been highly encouraged, federal requirements were needed to drive certain industries to participate. The use of procurement requirements has led to increased participation, as the opportunity to win federal contracts has incentivized incident-reporting practice adoption.

Industries such as finance, health, and critical infrastructure have cyber-incident-reporting requirements as part of law, but much of private industry is merely encouraged to submit incident reports voluntarily. One way the federal government has incentivized incident reporting is through its purchasing power by making it a requirement of earning federal contracts. DoD procurement rules require suppliers to report any vulnerability or compromise that impacts the contractor information system, the data residing in it, or the ability of the contractor to complete contractual requirements.[42] Adopting and implementing incident-reporting practices is then a necessity for any company looking to bid on a DoD contract.

The ability to share incident reports is critical to cybersecurity risk management practices in the federal government and across the private sector. President Biden's

2021 Executive Order on Improving the Nation's Cybersecurity directed the removal of barriers to sharing cyber threat and vulnerability information, and specifically called to identify contractual terms or restrictions that limit supplier sharing of incidents with federal agencies.[43] This directive was followed with the passage of the Cyber Incident Reporting for Critical Infrastructure Act of 2022, which institutes a requirement for federal agencies to share cyber incident reports with the U.S. Department of Homeland Security and the Cybersecurity and Infrastructure Security Agency.[44] Executive orders and legislation have highlighted the need for information sharing in mitigating cybersecurity risks and the value of federal procurement policies and contracting guidance to incentivize this practice.

Recommendation: Use contracting rules to incentivize and, when necessary, compel government suppliers to comply with AI incident reporting and cross-agency sharing.

The key lessons in the use of federal procurement requirements related to cybersecurity incident reporting are that contractual requirements provide incentives for getting industry to participate in incident reporting and that contractual terms must be written to enable the sharing of incident information. AI incidents may differ from cyber incidents, but the growth of common datasets, architectures, algorithms, and pretrained models has expanded the reach of harms and vulnerabilities.

Cyber incidents are typically actions by a malicious actor that undermine the confidentiality, availability, or integrity of information or an information system. AI incidents are different, in that they can include the operations of the model for its intended use, as well as manipulations by malicious actors. AI incidents can include unintended harm due to issues of bias and fairness, safety, privacy and rights, ethical concerns, and legal and regulatory issues. AI models can also be susceptible to external inputs and adversarial AI attacks, which aim to have the model learn, do, or display something incorrect. Though the two types differ, AI and cyber incident reports are necessary both to inform risk and to facilitate mitigation efforts.

The AI systems of today are rarely developed in the same way as the siloed models of the past. They are now often created using common datasets, frameworks, and pretrained models. The expansion of open-source frameworks and models exacerbates the commonality of systems.[45] As common AI system architecture components proliferate across industries, incidents that occur with one AI system may provide awareness of a widespread vulnerability or risk.

While the NIST AI RMF contains guidance on the development of incident reports, the president's executive order and the corresponding OMB guidance do not provide direction on the incident reporting of AI harms. The RMF, however, does identify that

some AI sectors have established harm reporting, disclosure, and documentation practices that could align with the established reporting requirements for cyber incidents. To effectively create AI incident-reporting constructs, the federal government should learn from the cybersecurity experience and establish as much commonality and centralization in the incident-reporting processes as possible. The reporting should then be mandated and enforced through federal procurement rules to incentivize incident reporting among the federal supplier base. Additionally, these rules should ensure that the contractual terms, and the corresponding reports, are structured in a way to encourage the most sharing across federal agencies.

Summary of Recommendations for AI Risk Management

Federal procurement rules can be an effective tool to ensure that suppliers comply with risk management practices in key technologies. The current frameworks and rules being developed by the federal government to manage the risks of AI are similar to those in cybersecurity. As Congress and federal agencies look to procurement rules to mandate AI standards, they should embrace the lessons learned from cyber procurement requirements by considering the following recommendations:

1. Develop standards to assess the level of risk and potential impacts of AI systems. Establish categories to differentiate the levels of risk AI systems pose and develop the appropriate risk management practices required for each category.

2. Base the level of requirements verification on the overall risk of the system. Compliance audits are costly, and therefore the federal government should utilize risk categories to determine which systems require compliance auditing.

3. Mandate and provide training on AI risk management standards for the federal acquisition workforce.

4. Leverage third-party auditors to support assessments of supplier compliance with AI risk management standards. This would solve labor limitations and skills gaps in the federal workforce, but it is important that final approval decisions rest with the government. Establish an AI standards center of excellence to provide government oversight and support compliance assessments.

5. Use contracting rules to incentivize and, when necessary, compel government suppliers to comply with AI incident reporting and cross-agency sharing.

Conclusion

While managing the risks of information system cybersecurity is different from managing the risks of an AI system, the challenges in ensuring compliance by federal suppliers are similar. Because of this, any AI procurement policy should be developed with insights derived from the experience of cybersecurity implementations to guarantee that rule developers are not repeating the mistakes of the past. As the federal government charts a course to govern AI risk management compliance, it should make sure that it has plans to overcome these challenges before codifying new rules.

Author

Matthew Schoemaker is a Major in the U.S. Air Force and is an Air Force Fellow at CSET.

Acknowledgments

I want to express my gratitude to Will Roberts of ASI Government and Jacob Horne of Summit 7 for serving as outside reviewers. I am also grateful for the generous support from CSET's John Bansemer, Heather Frase, Emelia Probasco, Owen Daniels, Mina Narayanan, Jenny Jun, and Michael O'Connor during the review and publication process. Not least, I would like to thank my wife for her faithful dedication and support.

Disclaimer

The conclusions and opinions expressed in this research paper are those of the author and do not necessarily reflect the official policy or position of the U.S. government, the U.S. Department of Defense, the Department of the Air Force, or Air University.

Document Identifier: doi: 10.51593/20240021

Endnotes

[1] Exec. Order No. 14110, 88 FR 75191 (2023), www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/.

[2] Shalanda Young, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," memorandum, Office of Management and Budget, March 28, 2024, www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf.

[3] NIST, *Initial Summary Analysis of Responses to the Request for Information (RFI) Evaluating and Improving Cybersecurity Resources: The Cybersecurity Framework and Cybersecurity Supply Chain Risk Management* (Washington, DC: Department of Commerce, 2022), www.nist.gov/system/files/documents/2022/06/03/NIST-Cybersecurity-RFI-Summary-Analysis-Final.pdf.

[4] Exec. Order No. 14110.

[5] Young, "Advancing Governance, Innovation, and Risk Management."

[6] Ted W. Lieu, Zoe Lofgren, Haley Stevens, letter to Shalanda Young on NIST AI RMF, July 20, 2023, https://lieu.house.gov/sites/evo-subsites/lieu.house.gov/files/evo-media-document/letter-to-omb-on-nist-ai-rmf-final.pdf; Rayid Ghani, "Governing AI through Acquisition and Procurement," Testimony to the Senate Committee on Homeland Security and Governmental Affairs, 118th Congress, September 14, 2023, www.hsgac.senate.gov/library/files/testimony-ghani-2023-09-14/testimony-ghani-2023-09-14/; Brad Smith, "Developing and Deploying AI Responsibly: Elements of an Effective Legislative Framework to Regulate AI," Testimony to the Senate Judiciary Committee, Subcommittee on Privacy, Technology, and the Law, 118th Congress, September 12, 2023, www.judiciary.senate.gov/download/2023-09-12-pm-testimony-smith?download=1; Justin Hendrix, "Transcript: Senate Hearing on the Need for Transparency in Artificial Intelligence," *Tech Policy Press*, September 13, 2023, https://techpolicy.press/transcript-senate-hearing-on-the-need-for-transparency-in-artificial-intelligence/; Lynne Parker, "Artificial Intelligence in Government," Testimony to the Senate Committee on Homeland Security and Governmental Affairs, 118th Congress, May 16, 2023, www.hsgac.senate.gov/hearings/artificial-intelligence-in-government/testimony-parker-2023-05-16-2/.

[7] White House, "FACT SHEET: Biden-Harris Administration Secures Voluntary Commitments from Eight Additional Artificial Intelligence Companies to Manage the Risks Posed by AI," September 12, 2023, www.whitehouse.gov/briefing-room/statements-releases/2023/09/12/fact-sheet-biden-harris-

administration-secures-voluntary-commitments-from-eight-additional-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai/.

[8] National Artificial Intelligence Initiative Act 2020, Pub. L. No. 116–283, 134 Stat. 3415 (2021), www.congress.gov/116/plaws/publ283/PLAW-116publ283.pdf.

[9] Artificial Intelligence Training for the Acquisition Workforce Act, Pub. L. No. 117–207, 136 Stat. 2238 (2022), www.congress.gov/117/plaws/publ207/PLAW-117publ207.pdf.

[10] Exec. Order No. 14110.

[11] Young, "Advancing Governance, Innovation, and Risk Management."

[12] Federal Artificial Intelligence Risk Management Act of 2023, S.3205, 118th Congress (2023), www.moran.senate.gov/public/_cache/files/f/2/f28eff88-9d3d-475d-bcb0-df22da1532e6/8DA77388E7C4B5FE471EDB8CB9840EF5.bag23e44.pdf.

[13] Justin Siken, "Record $765B in Federal Contracts Awarded in 2023," HigherGov, January 17, 2024, www.highergov.com/reports/765b-federal-gov-contract-awards-2023/; Nestor Maslej et al., *Artificial Intelligence Index Report 2023* (Stanford, CA: Stanford University Human-Centered Artificial Intelligence, 2023), https://aiindex.stanford.edu/wp-content/uploads/2023/04/HAI_AI-Index-Report_2023.pdf.

[14] "As AI development accelerates, private industry has yet to standardize practices for evaluating AI systems for risk, trustworthiness, and responsibility. Through federal procurement policy, the government has a unique opportunity to shape standards and frameworks for development and deployment of these technologies across the private sector more broadly" (Gary Peters, "Chairman Peters Opening Statement as Prepared for Delivery Full Committee Hearing: Governing AI Through Acquisition and Procurement," Senate Committee on Homeland Security and Government Affairs, September 14, 2023, www.hsgac.senate.gov/library/files/opening-statement-peters-2023-09-14/opening-statement-peters-2023-09-14/).

[15] NIST, *Initial Summary Analysis*.

[16] "Building on the model of existing rules that require federal contractors to adopt strong cybersecurity practices, Congress could likewise encourage industry adoption of standards based on the AI RMF by requiring federal contractors to self-attest, as a condition of bidding on federal contracts" (Smith, "Developing and Deploying AI Responsibly"); "Notably, NIST's impressive work on cybersecurity standards provides us with a clear precedent. We already task federal agencies with following NIST's standards and guidelines on cybersecurity. . . . With AI, the Administration can follow a similar path and

ensure agencies and vendors follow NIST's standards and guidance as a baseline to mitigate the risks and possible harms of the technology" (Lieu, Lofgren, and Stevens, letter to Young).

[17] NIST (Computer Security Division), *Minimum Security Requirements for Federal Information and Information Systems*, FIPS Pub. 200, March 2006, https://doi.org/10.6028/NIST.FIPS.200.

[18] Young, "Advancing Governance, Innovation, and Risk Management."

[19] NIST, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (Washington, DC: Department of Commerce, 2023), https://doi.org/10.6028/NIST.AI.100-1.

[20] European Commission, "AI Act," Directorate-General for Communications Networks, Content, and Technology, last updated May 6, 2024, https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai.

[21] Ellen Lord, "Assessing Contractor Implementation of Cybersecurity Requirements," memorandum, Under Secretary of Defense (Acquisition and Sustainment), November 14, 2019, www.acq.osd.mil/asda/dpc/cp/cyber/docs/safeguarding/14Nov2019-USD(A&S)-Memo-Assessing-Contractor-Implementation-of-Cybersecurity-Requirements.pdf.

[22] Office of Inspector General, "Audit of Protection of DoD Controlled Unclassified Information on Contractor-Owned Networks and Systems DODIG-2019-105," DoD, July 23, 2019, www.dodig.mil/reports.html/Article/1916036/audit-of-protection-of-dod-controlled-unclassified-information-on-contractor-ow/.

[23] Lord, "Assessing Contractor Implementation of Cybersecurity Requirements."

[24] NIST, *NIST SP 800-171 DoD Assessment Methodology, Version 1.2.1* (Washington, DC: Department of Commerce, 2020), www.acq.osd.mil/asda/dpc/cp/cyber/docs/safeguarding/NIST-SP-800-171-Assessment-Methodology-Version-1.2.1-6.24.2020.pdf.

[25] Chief Information Officer, "About CMMC," DoD, accessed November 2, 2023, https://dodcio.defense.gov/CMMC/About/.

[26] Chief Information Officer, "About CMMC."

[27] DoD, Cybersecurity Maturity Model Certification Program, proposed rule, 88 FR 89058, December 26, 2023, www.federalregister.gov/documents/2023/12/26/2023-27280/cybersecurity-maturity-model-certification-cmmc-program.

[28] "DoD", Cybersecurity Maturity Model Certification Program.

[29] "However, the proposed rule requires significant adjustments to balance security requirements with implementation costs" (Rachel McCaffrey and Michael Seeds, "The Costs and Scope of CMMC 2.0," National Defense Industrial Association, January 26, 2024, www.nationaldefensemagazine.org/articles/2024/1/26/ndia-policy-points-the-costs-and-scope-of-cmmc-20.

[30] Defense Contract Management Agency, "Top 'Other Than Satisfied' Requirements from DIBCAC High Assessments," PowerPoint presentation, December 2022, www.dcma.mil/Portals/31/Documents/DIBCAC/DIBCAC_Top_OTS_Reqts.pptx.

[31] Eric Noonan and Carl Herberger, "Defenseless: A Statistical Report on the State of Cybersecurity Maturity across the Defense Industrial Base (DIB)," webinar, CyberSheath, 2022, https://cybersheath.com/resources/webinars/defenseless/.

[32] Young, "Advancing Governance, Innovation, and Risk Management."

[33] Office of Inspector General, "Audit of Protection of DoD Controlled Unclassified Information."

[34] DoD Instruction 5200.48 "Controlled Unclassified Information", Office of the Under Secretary of Defense for Intelligence and Security, March 6, 2020, https://www.dodcui.mil/Portals/109/Documents/Policy%20Docs/DoDI%205200.48%20CUI.pdf.

[35] Under Secretary of Defense (Intelligence and Security), *Audit of the DoD's Implementation and Oversight of the Controlled Unclassified Information Program*, DODI 5200.48, June 1, 2023, https://media.defense.gov/2023/Jun/01/2003234002/-1/-1/1/DODIG-2023-078.PDF.

[36] Artificial Intelligence Training for the Acquisition Workforce Act; Exec. Order No. 14110.

[37] Frank Kendall, "Cyber Maturity Model Certification: An Idea Whose Time Has Not Come and Never May," *Forbes*, April 29, 2020, www.forbes.com/sites/frankkendall/2020/04/29/cyber-security-maturity-model-certificationan-idea-whose-time-has-not-come-and-never-may/?sh=7c0ba03a3bf2.

[38] Chief Information Officer, "CMMC Frequently Asked Questions," DoD, accessed November 27, 2020, https://dodcio.defense.gov/CMMC/FAQ/. The CMMC proposed rule estimates that 143,201 of 221,286 federal suppliers will only need to complete self-assessments and no longer require government or C3PAO assessment (DoD, Cybersecurity Maturity Model Certification Program).

[39] FedRAMP, "About FedRAMP Marketplace," FedRAMP Program Management Office, accessed January 8, 2024, https://marketplace.fedramp.gov/products?status=authorized.

[40] Ellen Lord, "Strategically Implementing Cybersecurity Clauses," memorandum, Under Secretary of Defense (Acquisition and Sustainment), February 5, 2019, www.acq.osd.mil/asda/dpc/cp/cyber/docs/safeguarding/USA000261-19-USD-Signed-TAB-A.pdf.

[41] Patrick Tremblay, "Building a Cybersecurity Assessment Capability," Defense Contract Management Agency Public Affairs, June 24, 2019, www.dcma.mil/News/Article-View/Article/1885182/building-a-cybersecurity-assessment-capability/.

[42] Defense Federal Acquisition Regulations System, Safeguarding Covered Defense Information and Cyber Incident Reporting, DFARS 252.204-7012, accessed December 5, 2023, www.acquisition.gov/dfars/252.204-7012-safeguarding-covered-defense-information-and-cyber-incident-reporting.

[43] Exec. Order No. 14028, 86 FR 26633 (2021), www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/.

[44] Cybersecurity and Infrastructure Security Agency, "Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA) Fact Sheet," accessed December 5, 2023, www.cisa.gov/sites/default/files/2023-01/CIRCIA_07.21.2022_Factsheet_FINAL_508%20c.pdf.

[45] Examples of common data, frameworks, and models include datasets like Common Crawl, ImageNet, and RefinedWeb; common AI frameworks of PyTorch and Tensorflow; and pretrained models like BERT and LLAMA.