

American University Washington College of Law

Digital Commons @ American University Washington College of Law

Joint PIJIP/TLS Research Paper Series

2021

The Integration of Artificial Intelligence in the Intelligence Community: Necessary Steps to Scale Efforts and Speed Progress

Corin R. Stone

Follow this and additional works at: <https://digitalcommons.wcl.american.edu/research>



Part of the [Intellectual Property Law Commons](#), and the [Science and Technology Law Commons](#)

The Integration Of Artificial Intelligence In The Intelligence Community

NECESSARY STEPS TO SCALE
EFFORTS AND SPEED PROGRESS

Corin R. Stone

Scholar-in-Residence and Adjunct Professor
Tech, Law & Security Program
Washington College of Law
American University

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	3
INTRODUCTION	4
PART I: CULTURE	8
<i>Creating an IC Innovation Culture</i>	10
<i>A National Security Ecosystem and Ethos</i>	13
PART II: OVERCOMING THE HURDLES	15
<i>Budget and Acquisition – Intertwined, Inflexible, and Complex</i>	15
FLEXIBILITY & SPEED: ISSUES	16
FLEXIBILITY & SPEED: SOLUTIONS	18
SOOs v. SOWs	18
Additional Authorities	20
Unclassified Sandbox	20
Single Appropriation	21
No-Year Funds	22
AI Technology Fund	23
ICWERX	23
COMPLEXITY: ISSUES	26
COMPLEXITY: SOLUTIONS	27
De-mystify Current Authorities	27
<i>Develop Partnerships</i>	27
<i>Prioritize Training & Education</i>	28
<i>Hire Expertise</i>	28
<i>Finalize Reference Book</i>	29
Incentivize Behavior	29
<i>Risk – A Scary Prospect</i>	32
ELEMENTS OF AI RISK FRAMEWORK	34
EXTERNAL STAKEHOLDERS	37
<i>Meaningful Congressional Oversight – A Critical Component</i>	38
ADAPTIVE OVERSIGHT	38
TRUST, TRANSPARENCY, AND PARTNERSHIP	41
PART III: SUMMARY OF ACTIONS	46
APPENDIX A: RISK FRAMEWORK INTRODUCTION	50
APPENDIX B: PREPARING THE FRAMEWORK	51
APPENDIX C: USING THE FRAMEWORK	54
APPENDIX D: BALANCING ACTION VERSUS INACTION	61
APPENDIX E: EXAMPLES	62

ACKNOWLEDGEMENTS

The author would like to thank the entire [Tech, Law & Security Program](#) at the American University Washington College of Law for their input, advice, and support; the many people inside and outside of government who answered questions, reviewed drafts, and offered feedback; and her tireless research assistants Meghan Anand, Sydney Jackson, John Jankosky, and Justin Singh.

ABOUT THE AUTHOR

[Corin Stone](#) is on loan from the Office of the Director of National Intelligence as a Scholar-in-Residence in the Tech, Law & Security Program at American University's Washington College of Law. She has held numerous senior roles in the U.S. Intelligence Community, including Deputy Director of National Intelligence for Strategy & Engagement, Executive Director of the National Security Agency, and Principal Deputy General Counsel for the ODNI. All opinions expressed herein are those of the author and do not reflect the official position or views of the ODNI or any other U.S. Government agency.

INTRODUCTION

Artificial Intelligence (AI)¹ is at the cutting edge of the technological revolution the world is experiencing. There is almost no industry left untouched by today's technological advances, including national security, which is becoming impossibly intertwined with emerging technologies like AI.² For example, China, is leveraging academic and private sector innovation and cutting-edge technology to directly benefit its national security and defense interests at record speed.³ But China is not the only country taking advantage of technology for national security purposes. A few decades ago sophisticated technology like AI was available only to a handful of wealthy governments, but now a range of advanced technology is available to anyone with an internet connection. It is ubiquitous and often easy to acquire, enabling smaller foreign governments and non-governmental actors alike to take advantage of it.⁴ Moreover, technology is "being invented, used, spread, and then discarded at ever increasing speeds around the world."⁵

It is already clear that how a government develops and employs emerging technology like AI raises a number of unresolved questions, including the ethics of doing so, potential bias of data and systems, and the legitimacy and reliability of AI outputs.⁶ These questions go to the very core of our values as a democratic nation and cause concern about the US Government's (USG) ability to use these tools. And yet, the USG cannot ignore AI retain its advantage; in an environment of great power competition and global access to technology, the USG must employ state-of-the-art AI tools to both respond effectively to emerging threats and to continue to lead on the world stage.⁷

The US Intelligence Community (IC), with its mission to outmatch US adversaries and discover actively concealed secrets to provide greater decision space to US policymakers, is no exception. Without the benefit of cutting-edge AI tools, being used by our adversaries all over the world, the IC will become ineffective and obsolete, leaving the nation vulnerable and exposed.⁸ As the National Security Commission on AI (NSCAI) noted in its Final Report, "[w]e must adopt AI to change the way we defend America, deter adversaries, use intelligence to make sense of the world, and fight and win wars. The men and women who protect the United States must be able to leverage the AI

1 The term "AI" is used throughout this report but the findings and recommendations are relevant to emerging technology more broadly, as well.

2 *Global Trends 2040: A More Contested World*, Nat'l Intelligence Council 1, 54, 64 (Mar. 2021), https://www.dni.gov/files/ODNI/documents/assessments/GlobalTrends_2040.pdf. See Lloyd J. Austin III, *Opinion: The Pentagon Must Prepare for a Much Bigger Theater of War*, Wash. Post (May 5, 2021), https://www.washingtonpost.com/opinions/lloyd-austin-us-deter-threat-war/2021/05/05/bed8af58-add9-11eb-b476-c3b287e52a01_story.html and Brendan McCord & Zoe A.Y. Weinberg, "How the NSC can better tackle emerging technology threats," Brookings (Feb. 1, 2021), <https://www.brookings.edu/techstream/how-the-nsc-can-better-tackle-emerging-technology-threats>.

3 Adam Segal, "Innovation and National Security: Keeping Our Edge," Independent Task Force Report No. 77 (Council on Foreign Relations, Sept. 2019), <https://www.cfr.org/report/keeping-our-edge> (describing the civil-military fusion pillar of Chinese military modernization in an effort to bolster the country's innovation system for advanced multiuse technologies in aviation, aerospace, and information technology); Michael Brown, Eric Chewning, & Pavneet Singh, "Preparing the United States for the Superpower Marathon with China," Global China (Brookings Institution, Apr. 2020), https://www.brookings.edu/wp-content/uploads/2020/04/FP_20200427_superpower_marathon_brown_chewning_singh.pdf (describing China's embrace of civil and military technology fusion).

4 T.X. Hammes, *Technology Converges; Non-State Actors Benefits*, Hoover Inst. (Feb. 25, 2019), <https://www.hoover.org/research/technology-converges-non-state-actors-benefit#:~:text=Further%2C%20if%20one%20studies%20their%20use%20of%20technology%2C,They%20seemed%20to%20do%20so%20for%20two%20reasons;See%20How%20Have%20Governments%20Changed%20with%20Technological%20Advances?>, Ctr. Public Impact (Feb. 28, 2017), <https://www.centreforpublicimpact.org/insights/governments-changed-technological-advances>.

5 *Global Trends 2040*, *supra* note 2, at 7.

6 *Id.* at 59.

7 *Final Report*, Nat'l Sec. Comm'n Artificial Intelligence 1, 24 (Mar. 1, 2021), https://assets.fole.com/eu-west-2/uploads-7e3kk3/48187/nscai_full_report_digital.04d6b124173c.pdf; see also Michèle A. Flournoy et al., *Building Trust through Testing: Adapting DOD's Test & Evaluation, Validation & Verification (TEVV) Enterprise for Machine Learning Systems, Including Deep Learning Systems*, Ctr. Sec. & Emerging Tech. 1, 4 (Oct. 2020), <https://cset.georgetown.edu/wp-content/uploads/Building-Trust-Through-Testing.pdf> ("[w]ithout urgent reforms and prioritized investment . . . DOD . . . will lose the opportunity to take advantage of new private sector developments, while allowing other nations without such standards to adopt the latest innovations").

8 *Global Trends 2040*, *supra* note 5, at 58-9.

and associated technologies that can help them accomplish their missions as quickly and safely as possible.”⁹

For these reasons, government commissions, legislators, private sector studies, and international task forces, among others, have put forth hundreds of thoughtful recommendations as to how the United States can best take advantage of the opportunities and protect against the risks brought by AI. Indeed, the NSCAI just finished three years of work resulting in sixteen chapters and many appendices full of recommendations for the US Government. This work is a critical starting point and the US Government must continue to take advantage of AI in appropriate ways. The IC cannot wait to solve every issue and buy down all risk before it begins to leverage these capabilities. The AI train has left the station and the US Government has no option but to build and run the train at the same time.

This report presumes that questions about ethics, privacy, transparency, and security are critical and must be considered before the IC moves forward with AI at scale. And the IC has taken initial steps to begin that conversation.¹⁰ The focus of this report recognizes that even as those critical issues continue to take shape, there are existing core areas where the IC still must fundamentally change the paradigm if it is to take full advantage of AI: the basic mechanics of acquiring and developing it.¹¹

The IC has impressive scientists, engineers, mathematicians, and others developing incredible technological breakthroughs in various organizations. The Intelligence Advanced Research Projects Activity (IARPA)¹² and other IC research organizations focus on basic and applied, high-risk, high-reward research with innovative solutions that have low or unknown likelihood of success. However, IARPA looks for solutions in the 3-5 year timeframe, so these solutions will neither materialize nor scale quickly enough for immediate technological needs.

In-Q-Tel, a government-sponsored but private venture capital company created in the late 1990s to connect the CIA to Silicon Valley, is focused on making strategic investments to shape and accelerate private sector innovation – another critical need of the IC. It does this by investing in visionary start-up companies and connecting their cutting-edge technologies to specific IC mission needs to influence the direction of a company’s work.¹³ However, by design, In-Q-Tel is independent from the IC and makes its own decisions about investment-worthy technology. In-Q-Tel is not beholden to IC-wide priorities nor does it take direction from the IC. In addition, In-Q-Tel does not facilitate direct IC engagement with the private sector and does not produce immediate solutions. So even with the great work of In-Q-Tel, IARPA, and other IC research organizations there is a missing piece for the IC: the ability to easily engage directly with private sector entities, FFRDCs, academia and others to quickly take advantage of their ongoing and evolving AI efforts.

This gap is compounded by the fact that the IC is not currently designed for technical speed or agility but is trying to engage with companies and technologies that require it. The IC’s budget and acquisition processes, modeled on DOD processes, are Cold War relics built around outdated

⁹ *Final Report*, *supra* note 7, at 21.

¹⁰ *Principles of Artificial Intelligence for the Intelligence Community*, Office Dir. Nat’l Intelligence, https://www.dni.gov/files/ODNI/documents/Principles_of_AI_Ethics_for_the_Intelligence_Community.pdf.

¹¹ Segal, *supra* note 3, (“The [DOD] and the [IC] will fall behind potential adversaries if they do not rapidly access and deploy technologies developed in the private sector.... A persistent cultural divide between the technology and policymaking communities threatens national security by making it more difficult for the [DOD] and [IC] to acquire and adopt advanced technologies from the private sector and to draw on technical talent.”).

¹² *IARPA*, Office Dir. Nat’l Intelligence, <https://www.iarpa.gov>.

¹³ *About, IN-Q-TEL*, <https://www.iqt.org/about-iqt/>.

beliefs that the threats to our nation are stable and predictable, and that the government is primarily responsible for the development of cutting-edge technology.¹⁴ Neither of these beliefs remains true and so the outdated acquisition and budget frameworks within which the IC still operates have become arduous and incompatible with what this moment demands.

More specifically, the USG's budget processes are neither flexible nor fast enough to adapt to evolving requirements for AI.¹⁵ For example, federal agencies must develop budget requests several years in advance of execution, and Congress significantly limits the agencies' abilities to use funds and adjust for changes in the year of execution. Resources often expire within a year if not already spent, and different money must be used for each phase of the acquisition lifecycle.

In addition, the USG's acquisition processes are complex, rigid, and slow, favoring bigger companies with sufficient resources to navigate the bureaucracy, rather than small, innovative start-ups that often create new, cutting-edge technologies and tools, like AI. And the complexity of the acquisition processes confounds even those within the USG, with many acquisition professionals unfamiliar with or uncomfortable using what flexibilities they do have.

The interrelated problems of antiquated budgeting and cumbersome acquisition processes are compounded by two additional areas that the USG must re-imagine: risk and oversight. Emerging technologies like AI are not always initially successful and almost always bring a risk of failure. However, in the USG, risk is not rewarded; failure, even with noble intent and reasonable impact, often results in severe negative ramifications for those officers associated with it. There is no agreed-upon risk assessment framework to help ensure an understanding of or support for risky actions and potential failures. There is a vital difference between an innovative project that fails and a failure to innovate. The former teaches us something we did not know, while the latter is a national security risk. Without a considered and recognized approach to risk and failure, the IC workforce will not pursue cutting-edge technologies like AI.

And both the USG and Congress must radically re-think oversight processes. Congressional oversight is critical and non-negotiable. However, oversight needs adjustments so that it is focused on the right goals and at the right level. Traditional oversight revolves around predictability, certainty, and clarity of USG actions, which is largely incompatible with the development of AI. Oversight can become too tactical, creating unnecessarily cumbersome engagements that are unsatisfying to both the executive and legislative branches, and often lead to discord, distraction, and delay.

To date, smart people with dogged patience and perseverance have worked to clear these hurdles. At the leadership level, many have and continue to push on the unnecessary layers of bureaucracy that bring even the strongest officers to a screeching halt. At the implementation level, resilient and creative officers discover workarounds and exceptions, brute-forcing new technologies and innovations into the US Government. However, *ad hoc* success achieved by fighting through obstacles or barriers, whether or not they are well-intentioned, is not a good indicator of the overall health of agency organization or culture; that organization and culture must instead cultivate tech

¹⁴ Dan Patt & William Greenwalt, *Competing in Time: Ensuring Capability Advantage and Mission Success through Adaptable Resource Allocation*, Hudson (Feb. 25, 2021), <https://www.hudson.org/research/16717-competing-in-time-ensuring-capability-advantage-and-mission-success-through-adaptable-resource-allocation>; Robert F. Hale, *Financing the Fight, A History and Assessment of Department of Defense Budget Formulation Processes*, Brookings (Apr. 2021), <https://www.brookings.edu/research/financing-the-fight-a-history-and-assessment-of-department-of-defense-budget-formulation-processes/>.

¹⁵ *Interim Report*, Nat'l Sec. Comm'n Artificial Intelligence 1, 28-31 (Nov. 2019), https://www.nscai.gov/wp-content/uploads/2021/01/NSCAI-Interim-Report-for-Congress_201911.pdf.

talent, streamline process, and encourage calculated risk-taking to ensure widespread success on technical matters. Our AI needs are too great and time is running out – we no longer have the luxury of relying on a handful of resourceful officers to help one or two organizations take advantage of emerging technology. We must provide the support and frameworks leaders need to quickly push change through the IC on a macro scale.

While there is no silver bullet, this report fuses existing and new recommendations to prioritize actions that will drive faster, more effective transition of cutting-edge AI into the IC. This report is aimed at helping and garnering more support for executive and legislative branch leaders who are already tackling these issues. In furtherance of that goal, this report proposes implementation steps, where appropriate.

Part I of this report addresses an overarching issue that permeates everything else: culture. Part II tackles each of the areas identified above, in turn. Part III provides a summary of the proposed actions that would deal a direct blow to these hurdles. The appendices set forth a draft IC Risk Framework with examples as a starting point for the IC. Taken together, this report prioritizes the steps to much-needed relief that will allow the IC to quickly embrace AI, as those hundreds of thoughtful recommendations intend.

PART I: CULTURE

To successfully tackle change in the areas of budget, acquisition, risk, and oversight, the IC must simultaneously prioritize an issue more intangible and nebulous – its own culture. Culture is the ethos of an organization – the beliefs, behaviors, values, and characteristics of a group that are learned and shared over many years.¹⁶ In the IC, there are several predominant cross-cutting cultures,¹⁷ all of which flow from the mission of the IC – protecting the women and men on the front lines, defending U.S. national security, and delivering insights to policymakers at the speed of decision. This mission is a powerful and unifying force that leads to important IC values and behaviors.

Intelligence operations – uncovering foreign secrets and protecting assets, for example – are inherently risky; they very often put people in harm's way. If there is a leak of information related to an operation – if the people involved, or the location or target of an operation are exposed – not only might the mission fail to collect the desired information, but someone's life could be in jeopardy. The extreme consequences of leaks are well-understood, thanks to notorious spies like Robert Hanssen and inside leakers like Edward Snowden, but significant damage can also flow from what seem like merely small mistakes. If someone fails to make a connection between relevant information or forgets to check a database of known terrorists, for example, the results can be just as disastrous. Thus, the IC's high-stakes operations drive an enormous emphasis on security, preparation, and tradecraft, all of which help mitigate operational risk.

This same spirit manifests in “enabling” activities, like budget, acquisition, or cybersecurity, through a focus on certainty of action and predictability of results. Enabling activities are somewhat removed from the “pointy end of the spear” but are no less critical to the ultimate success of the mission. Proper funding and resources, the right capabilities, skilled officers, legal approval, and the many other support activities are integral to successful operations.

In the field, risks are unavoidable – operators cannot choose inaction to avoid those risks – given that risks are inherent in what they do, they must accept the reality that risks are inevitable, and they must learn to manage those risks to get the huge payoff of successful operations. So, the focus is on risk management rather than avoidance – what level of risk is acceptable for what level of intelligence gain?

Back home, where most enabling activities are handled, risks are not seen as inevitable – certainly not big ones. They are seen as avoidable, and subject to being minimized and mitigated. And some believe the best way to do that is by staying with tried-and-true standard operating procedures rather than experimenting with new approaches. Innovation is inherently risky, it can and will fail, especially early on. Innovation is not mandatory, it is entirely optional. Therefore, if the tendency is to avoid risks, in most cases innovation will be avoided.

In addition to this instinct, there are a variety of compounding issues that discourage creativity and innovative change. First, there are practical difficulties: change is hard, messy, and requires resources that most offices cannot spare. These concerns alone are big hurdles to clear. Second, innovative change means uncertainty – in execution, accountability, and success. And that uncertainty leads to the risk that projects may fail, resulting in loss of money, reputation, or even position. Thus,

¹⁶ *Organizational Culture*, gothamCulture, <https://gothamculture.com/what-is-organizational-culture-definition/>.

¹⁷ As opposed to agency-specific cultures, which also exist.

control, compliance, and trust are paramount and there is a strong aversion to things “not invented here.” Third, people are creatures of habit, and the IC is no different. It is easier, more comfortable, and safer to stick with what has been done in the past – even if not wholly successful – and to build our lives around the repetition of processes and activities.¹⁸ Indeed, we are taught as children to follow and embrace routines. Repetition and routine provide a level of safety and confidence that we are on a good path that involves little risk of failure. But this kind of approach also assumes our environment is static; that what we needed before is what we need now, and that the approach of the past is acceptable for the future, even if not always successful.

We do not live in a static environment, of course. And especially now, in the context of emerging technologies like AI, it is hardly an overstatement to say that our environment is changing almost daily. Not only the technology itself changes, but how we think about technology, how we develop it, how we fund it, how we engage with it, how it affects every aspect of our lives. The world around us is changing exponentially and the IC must consider the need to adapt its culture in significant ways – in terms of speed, agility, and willingness to take risks that will inevitably result in some level of failure.

However, even when it is needed, change is hard. And it is difficult, if not impossible, to create major cultural change in the absence of a compelling real-world problem.¹⁹ As leading change management expert John Kotter explains,²⁰ a sense of urgency helps others see the need for change and the importance of acting immediately. For example, prior to COVID-19, telework in the IC was virtually unheard of. Over the past 18 months, however, the IC has not only achieved telework for many IC roles, it proved to be so productive that many IC elements have now changed their standing policies to allow it in certain roles going forward. Using that crisis to create a blueprint for new long-term approaches was key to jump-starting a real transformation.

Another example occurred after the September 11, 2001 attacks (“9/11”). Prior to that event, the IC culture around sharing and protecting information was heavily weighted toward jealously guarding information within agency stovepipes. This was a result of many things – including the genuine need to protect sources and methods – but evolved in the extreme because it is easier to protect sensitive information if fewer people have access to it and, importantly, knowledge is power. After 9/11, though, the IC culture slowly started to evolve to one that acknowledged the need to connect the dots and proactively share information with critical partners, even while protecting it. The need to share information more freely within the IC had been documented for decades²¹ and executive branch policy had always set forth a requirement to allow information to go to those who had a “need to know”²² but it took the crisis of 9/11 to really break through the cultural barriers to information sharing.

18 See also *Sunk cost fallacy*, Behavioral Economics, <https://www.behavioraleconomics.com/resources/mini-encyclopedia-of-be/sunk-cost-fallacy/> (description of sunk cost fallacy).

19 Ignacio Crespo, et al., *The Evolution of Model Risk Management*, McKinsey & Co. 1, 4-5, 8 (Feb. 2017), <https://www.mckinsey.com/-/media/McKinsey/Business%20Functions/Risk/Our%20Insights/The%20evolution%20of%20model%20risk%20management/The-evolution-of-model-risk-management.pdf?shouldIndex=false>.

20 The 8-Step Process for Leading Change, Kotter Inc., <https://www.kotterinc.com/8-steps-process-for-leading-change/>.

21 Larry D. Thompson, *Intelligence Collection and Information Sharing within the United States*, Brookings (Dec. 8, 2003), <https://www.brookings.edu/testimonies/intelligence-collection-and-information-sharing-within-the-united-states/>.

22 See Richard A. Best Jr., Cong. Research Serv., RL33873, *Sharing Law Enforcement and Intelligence Information: The Congressional Role* (Feb. 13, 2007); see also Exec. Order No. 12958, 60 F.R. 48863 (Apr. 17, 1995).

Indeed, the IC has created an entire information technology enterprise architecture – groundbreaking for the US Government when it was started in 2011 – founded on the concept that the IC agencies must work together and share information more easily than in the past.²³ The IC Information Technology Enterprise (IC ITE) has been underway for a decade and is not yet complete, not solely due to its audacious technical goals but because it is breaking down cultural barriers, connecting the IC as never before, and pushing the IC together beyond its comfort level. Culture shifts slowly, and this one has had its detractors, but eventually over time there has become a general acceptance of and willingness to share information in ways that would previously have been unthinkable. The 9/11 crisis and subsequent terrorist attacks (and attempts) brought the responsibility to share information into immediate, sharp focus.

One of the IC's crises today is the ubiquity of and access to AI capabilities around the globe by state and non-state actors alike. Adversaries with access to AI capabilities go well beyond traditional nation states to non-state terrorist organizations, hackers, and criminals, among others, who use these tools for various goals, including to deny and deceive the IC from gaining insights.²⁴ If the IC does not modernize its approach to AI quickly, the U.S. will be severely disadvantaged in the international community. The IC must turn this crisis into an opportunity to embrace a more flexible and innovative culture that supports the transition of AI tools into the community at the speed of mission. As Winston Churchill admonished, the IC must not let this good crisis go to waste.

Creating an IC Innovation Culture

Culture is driven from the top through leadership actions, behaviors, priorities, and incentive structures that reverberate across every level of an organization.²⁵ In the IC, the top-down leadership required to prioritize and reward enterprise-wide innovation,²⁶ encourage reasonable risk-taking, support small failures, and ensure appropriate accountability has been fleeting in a way that has not allowed for a sustainable, pervasive IC-wide culture of innovation.²⁷

In 2018, the Director of National Intelligence (DNI) made IC-wide innovation a top priority by creating an ODNI innovation organization²⁸ to lead and support, among other things, cross-IC creativity and modernization – to help pioneers find each other, to encourage the IC to leverage each other's ideas, to provide advice to those hoping to contribute new ideas, and to train officers on innovation best practices.²⁹ This organization also led the IC-wide strategic initiative called

23 The IC IT Enterprise (IC ITE) is a strategy to move the IC from an agency-centric IT architecture to a common platform where the IC can easily and securely share technology, information, and resources. *IC IT Enterprise*, Office Dir. Nat'l Intelligence, <https://www.dni.gov/files/documents/IC%20ITE%20Fact%20Sheet.pdf>.

24 Paige Young, Artificial Intelligence: A Non-State Actor's New Best Friend, *OTH Journal* (May 1, 2019), <https://othjournal.com/2019/05/01/artificial-intelligence-a-non-state-actors-new-best-friend/>.

25 Boris Groysberg, et al., The Leader's Guide to Corporate Culture, *Harvard Business Review* (Feb. 2018), <https://hbr.org/2018/01/the-leaders-guide-to-corporate-culture>.

26 Innovation is often considered a buzzword devoid of meaning, resulting in a dismissive reaction. Defining innovation can help with this. MITRE suggests "novelty with impact." See Pete Modigliani, et al., FIVE BY FIVE: Five Disciplines and Five Strategic Initiatives for the Pentagon in the Digital Age, MITRE (Feb. 25, 2021), 1, 28, https://www.mitre.org/sites/default/files/publications/pr-20-03241-1-five-by-five-five-disciplines-and-five-strategic-initiatives-for-the-pentagon-in-the-digital-age_0.pdf. See also Steve Blank, www.steveblank.com, who has used the phrase "mission acceleration" in place of innovation to convey that innovation must lead to a valuable result.

27 Final Report, *supra* note, at 77 ("Despite pockets of imaginative reform and a few farsighted leaders, DoD remains locked in an Industrial Age mentality to which great-power conflict is seen as a contest of massed forces and monolithic platforms and systems.").

28 *Transformation & Innovation – Who We Are*, Office Dir. Nat'l Intelligence, https://www.dni.gov/index.php?option=com_content&view=article&id=2626:transformation-innovation-who-we-are&catid=318:transformation-innovation.

29 See MITRE, *supra* note 26, at 28 (explaining that leaders must ensure people know how to innovate. MITRE has an Innovation Toolkit that may be helpful).

Augmenting Intelligence using Machines (AIM),³⁰ which drove senior and subject-matter-expert coordination and collaboration across the IC on AI and machine learning activities to better align innovation, acquisition, and use of AI and emerging technology.

The AIM initiative was one of only six priorities very visibly embraced and supported by the directors of every IC element as imperative to the future of the IC. As a result, the AIM initiative had relative success in its first few years, bringing the Community together and ensuring better coordination across the many AI initiatives. However, the innovation office's more foundational task of creating a broad IC culture that embraces innovation foundered. Chronic under-resourcing of those foundational efforts signaled that they were not a priority and subsequent changes in ODNI leadership resulted in minimal progress over the course of two years. That office was ultimately disbanded,³¹ leaving IC innovators to continue to fend for themselves and signaling a lack of strategic support and leadership for IC-wide innovation.³²

There are glimmers of hope, however. Innovation remains one of the IC's stated values³³ and, despite the lack of coordinated approach, there are pockets of brilliant innovation across the IC because individual visionaries recognize that new ideas and technology can transform the world of intelligence. These enterprising and energetic individuals are expending enormous effort to find each other and work together, to create new pathways, build out novel ideas, and find new solutions to old problems. They are willing to fail and accept the consequences because they believe so strongly in the importance of their work.

To grow an IC-wide culture of innovation, IC leaders must harness this grassroots energy and enthusiasm and set the right vision and tone by affirmatively and publicly embracing outside-the-box thinking, articulating acceptable risk and expected failure rates, rewarding creativity even when a project fails,³⁴ and setting expectations for collaboration between innovators and practitioners in every area – from budget to acquisition to operations.³⁵ The DNI should immediately designate a senior leader with strong institutional credibility to drive innovation across the IC, starting with a 90-day action plan that takes advantage of the AI crisis to kick-start activities that will help shape the IC culture toward one that embraces and supports innovation more broadly. This action plan must create an incentive structure that rewards creativity and ensures officers are not punished for failures that result from reasonable risk taking; as long as officers are punished for failure, the culture will not change. Strong and steady leadership, clear prioritization, and a willingness to hold the organization accountable for achieving those goals are critical for creating and adapting a culture.

As in all organizations, some IC officers will embrace change quickly and be an active part of the solution. However, change can be hard and people with a long history in the IC can become attached to

30 *The AIM Initiative: A Strategy for Augmenting Intelligence Using Machines*, Office Dir. Nat'l Intelligence (Jan. 16, 2019), <https://www.dni.gov/index.php/newsroom/reports-publications/item/1940-the-aim-initiative-a-strategy-for-augmenting-intelligence-using-machines>.

31 *Organizational Chart*, Office Dir. Nat'l Intelligence, https://www.dni.gov/files/ODNI/documents/ODNI_Org_Chart_Final_For_Web_2020_1.pdf.

32 Final Report, *supra* note 7, at 128. ("Strategic initiatives succeed or fail at the tactical level . . .").

33 *Our Values*, Intel.gov, <https://www.intelligence.gov/mission/our-values/345-innovation>.

34 See MITRE, *supra* note 26, at 28.

35 See *Rightly Scaled, Carefully Open, Infinitely Agile: Reconfiguring to Win the Innovation Race in the Intelligence Community: Hearing before the Subcomm. On Strategic Technology and Advanced Research*, 116th Cong. 1, 22 (2020) (statement of Rep. Himes, Chairman & Rep. Stewart, Ranking Member) ("What makes for an innovative environment is subject to debate, but a variety of essential qualities are obvious: an open, collaborative culture, often between innovators and nearby academic and research institutions; a culture which embraces, rather than punishes risk-taking; an almost religious devotion to doing things differently—what economists call 'disruption' and what Mark Zuckerberg called 'breaking things.' These are not qualities that are readily embraced by the federal national security apparatus.").

existing processes and mechanisms, mistaking them for the ethos of the organization. And because major culture change cannot leave all existing processes untouched, it may provoke skepticism and pushback from those folks. But dragging them along or leaving them behind is not an option – officers with time and depth in the IC are critical; they make up the engine of the community and support and nurture employees throughout the IC. Actively engaging them as a part of the process, starting from a place of mutual pride in the IC and pursuing small changes that can lead to big impact, will help refresh the IC’s culture while preserving and championing its strengths, and pave the way for innovation and AI adoption at scale.³⁶

Similarly, executive and legislative branch overseers are comfortable and familiar with the decades-old processes that have worked well in many ways. They are unlikely to be enthusiastic about radical, comprehensive change without first seeing proofs of concept – evidence that the IC will continue to have success and overseers will continue to have the requisite insights and control despite whatever changes may take place. The IC must prove out a few impactful ideas to show the IC can handle increased flexibility and speed without losing the rigor and accountability expected and required.

The IC must quickly make inroads to adopt and inculcate an iterative, innovative culture by harnessing the real, current opportunities brought by AI. Because no matter how many brilliant minds come together to create excellent recommendations to take advantage of AI and promote innovation, the IC will not successfully implement them without addressing the institutional resistance to new ways of doing business that acts as a self-sustaining barrier between the IC and widespread adoption of AI.

ACTIONS: CULTURE OF INNOVATION

1. The DNI must clearly and publicly prioritize innovation. The DNI and IC leadership should:
 - a. Designate a senior leader to drive innovation across the IC as a top priority.
 - b. Set clear goals and expectations for IC innovation activities.
 - c. Articulate acceptable risk and expected failure rates.
 - d. Change incentive structures to reward creativity and innovation.
2. The DNI should require a 90-day action plan that takes advantage of the AI burning platform to drive a series of proposals and activities that will begin to turn the IC culture toward one that embraces and supports innovation broadly.
 - a. These proposals should articulate expectations for how innovation will be integrated across IC activities, rather than done in independent stovepipes.

³⁶ *Id.*

A National Security Ecosystem and Ethos

The IC and DOD are part of the same national security ecosystem – they have a common objective of safeguarding our nation’s security, they provide critical support to each other in carrying out their work, and they have significant overlapping and tightly coupled activities toward that end. Indeed, the DOD intelligence agencies³⁷ are part of both DOD and the IC for that very reason. The IC and DOD must and do work hand-in-hand when it comes to national security; it is key to ensuring that the USG operates as a unified, informed government when executing national security missions. This critical partnership has been recognized by many experts, resulting in recommendations to ensure continued tight coordination between the IC and DOD through, among other things, a “Steering Committee on Emerging Technology, tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.”³⁸ This and other recommendations, however, do not specifically focus on the need for better synchronization of IC and DOD authorities, a critical gap in how the IC and DOD complement each other.

Due in part to the significant overlap in missions and personnel mentioned above, the IC and DOD often have similar requirements in authorities. It is not uncommon that if DOD needs greater flexibility or new authority, the IC could also use some version of the same. To date, if an especially aware or imaginative officer recognizes the dual equities of DOD and the IC, they may propose similar legislation for the other. Likewise, the Office of Management and Budget (OMB), which coordinates interagency legislative proposals, may flag provisions of interest for other agencies. And then again there are times when the ODNI, which has both the mission and visibility to consider IC-wide activities and authorities, will propose changes to include more agencies or DOD in a program or authority. But this approach is ad-hoc and insufficient, as is evident by continuing instances of DOD and the IC seeking authorities that would be very helpful to each other but not recognizing or developing them in tandem. Neither IC nor DOD officers reflexively consider the entire national security ecosystem throughout the legislative proposal process. Agencies and officers sometimes coordinate on legislative requirements at the outset, but it is typically only after one goes to OMB or Congress that the other realizes they may have a stake in the outcome. Or even more often, it is not until years later.

This was exactly the case with a program called the Public-Private Talent Exchange, which allows employees to flow to and from the private sector and government more easily, sharing knowledge and expertise in both directions. DOD had this program in place for *three years* before a few enterprising folks in the ODNI found it and realized the IC could use it, too. Indeed, upon closer look, there was no reason why the IC should not have had it three years earlier. Now the IC is working to catch up with related policies and processes to get its first employees moving back and forth to and from the private sector, very much in need of this program³⁹ and very much behind the curve.

We also see this phenomenon with the current DOD activities related to budget and acquisition, spurred in earnest by the DOD Software Acquisition and Practices (SWAP) Study,⁴⁰ whose

37 The Defense Intelligence Agency, the National Security Agency, the National Geospatial-Intelligence Agency, and the National Reconnaissance Office, which is a joint-venture between DOD and the IC. In addition, the intelligence organizations within the military services are also part of the IC for this reason.

38 Final Report, *supra* note 7, at 289; see also Maintaining the Intelligence Edge: Reimagining and Reinventing Intelligence Through Innovation, Ctr. Strategic & Int’l Stud. 1, 28 (Jan. 2021), https://csis-website-prod.s3.amazonaws.com/s3fs-public/publication/210113_Intelligence_Edge.pdf.

39 Final Report, *supra* note 7, at 119.

40 Software Acquisition and Practices (SWAP) Study, Def. Innovation Bd. (Mar. 2019), <https://innovation.defense.gov/software/>.

recommendations for more flexible and agile approaches are extremely helpful and could greatly benefit the IC. Congress has passed many of the legislative proposals⁴¹ within the SWAP study for DOD, but in most instances has not included the IC in those authorities. When granting authorities that encourage innovation and remove bureaucratic impediments, Congress should consider granting those authorities simultaneously to DOD and the IC.

Going forward, it is imperative that the executive and legislative branches instinctively and systematically think about the national security community more holistically, up and down the chain of command. Despite the important and necessary distinctions between the IC and DOD that must continue, there are significant complementary and overlapping missions, personnel, agencies, and activities that very often require similar authorities and flexibilities.⁴² Ensuring tighter coordination on new authorities would ultimately lead to greater interoperability, higher productivity, increased efficiency, and more cohesive action across the IC and DOD. For this reason, when the IC or DOD receives a flexibility or pilot authority related to a common national security imperative, everyone – across the IC, DOD, OMB, and Congress – must ask themselves whether it is also something the other could use. The answer may not always be yes, but it often will be.

ACTIONS: NATIONAL SECURITY ECOSYSTEM

1. ODNI and DOD General Counsels and Legislative Affairs Chiefs should meet semi-annually to discuss proposed legislation that may be applicable to the broader IC and DOD.
 - a. They also should create a specific, formal step in the legislative process to ask whether the other might need similar legislation.
2. The ODNI and DOD should instruct all IC and DOD officers, with specific attention to legal and legislative officers, to proactively consider who else may need relevant authorities and propose adjustments, if appropriate. This includes ensuring subject matter experts are consulted to understand the objectives of the authorities and programs.
 - a. This is doubly important for the ODNI team, as they have a unique cross-IC optic and awareness and it is specifically their mission to look out for the entire community.
3. The NSCAI's Steering Committee on Emerging Technology should require annual briefings from the ODNI and DOD General Counsels and Legislative Affairs Directors on their efforts to better coordinate AI-related pilots, authorities, and proposals.

⁴¹ See also MITRE, *supra* note 26, at 12 (“Congress has been a great partner [for DOD] over the past five years to grant dozens of authorities in recent NDAs, helping to enable greater speed and flexibility as it delegated authorities and established new pathways.”).

⁴² Final Report, *supra* note 7, at 114 (“The IC must aggressively pursue automated interoperability with the DoD for intelligence operations conducted at machine speeds. . . . They should work together to create interoperable and sharable resources and tools . . . and should . . . [share] all AI-enabled capabilities whenever feasible.”).

PART II: OVERCOMING THE HURDLES

We are at a global inflection point in terms of the speed and depth of evolving technological capabilities that are radically changing the national security landscape. And the USG is no longer the only dominating force in creating and using new and innovative technologies; such technologies are often developed in the private sector and are available to allies and adversaries alike. As a result, when it comes to outpacing our adversaries, the IC has no choice but to partner smartly with the private sector to quickly take advantage of AI innovation and creativity that will lead to new opportunities and help manage ongoing challenges in the national security space.

But despite the urgency of this moment, the IC continues to be mired down in existing, foundational processes that make the acquisition and integration of AI into the IC a challenge. The IC's budget and acquisition processes are complex, inflexible, and slow, and there is a cultural aversion to risk that is compounded by the IC's oversight regimes. It is past time for a paradigm shift away from the IC's traditional, linear processes to more iterative and dynamic approaches that will speed and transform the way the IC purchases, incorporates, and manages the use of AI.⁴³

This section provides additional background on these challenges and proposes priority short and long-term actions to address them.⁴⁴

Budget and Acquisition – Intertwined, Inflexible, and Complex

The inflexibility and complexity of the budget and acquisition processes across the IC, DOD, and the U.S. government more broadly rise to the top of the issues that must be addressed before the IC can take full advantage of AI.⁴⁵ The government's budget and acquisition processes are distinct, but very much intertwined and unfortunately not well aligned.⁴⁶ Moreover, they often focus on describing the solutions we want to pursue, rather than the problems we must solve. With the speed technology evolves today, we do not always know the solution we want; we only know the problem that exists, and we need fast iteration and continuous learning to successfully solve it in time for operational relevance.

This paper will not rehash the significant scope of research and writing on DOD's problems regarding these entangled processes.⁴⁷ Importantly, however, many of the IC's processes are either the same as DOD's or modeled after them, and so just as DOD must reform its processes to support AI – by allowing for “speed, uncertainty, experimentation and continuous upgrades”⁴⁸ – so too must the IC.⁴⁹

43 Patt & Greenwalt, *supra* note 14, at 8.

44 During the research of this project, the Center for Strategic & International Studies produced a report entitled, *Maintaining the Intelligence Edge, Reimagining and Reinventing Intelligence through Innovation*. CSIS's report provides critically-needed focus on strategic issues facing the IC when it comes to innovation and beneficial recommendations for improvement that also touch upon areas in this report. We endorse that report's findings and conclusions and further prioritize, narrow, and refine specific steps in this paper. Ctr. Strategic & Int'l Stud., *supra* note 38, at 19.

45 J. Michael McQuade et al., *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, Def. Innovation Bd. 1, 31 (May 3, 2019), https://media.defense.gov/2019/Apr/30/2002124828/-1/-1/0/SOFTWAREISNEVERDONE_REFACTORINGTHEACQUISITIONCODEFORCOMPETITIVEADVANTAGE_FINAL.SWAP.REPORT.PDF. See Final Report, *supra* note 7, at 66.

46 Patt & Greenwalt, *supra* note 14, at 39 (describing the linear and serial process of requirements, budget, contracting, and acquisition, “which makes addressing requirements reform independent of budgetary reform intractable”); see also MITRE, *supra* note 26, at 9 (DOD leaders struggle to align and integrate acquisition, requirements, and budget).

47 Patt & Greenwalt, *supra* note 14; Hale, *supra* note 14; SWAP Study, *supra* note 40, at 19-21; see also MITRE, *supra* note 26.

48 Final Report, *supra* note 7, at 66.

49 Rightly Scaled, *supra* note 35.

FLEXIBILITY & SPEED: ISSUES

The Constitution vests Congress with the “power of the purse” – the ability to direct and oversee the collecting and spending of taxpayer dollars on behalf of the federal government. To carry out this duty and ensure the executive branch adheres to its direction, Congress crafted a variety of rules to ensure visibility, control, and accountability over the spending of government resources. In the executive branch, the Office of Management and Budget (OMB) also ensures wise and careful use of taxpayer dollars and, therefore, OMB is highly focused on efficiency and minimizing change to the pre-determined cost and schedule of projects.

Proper and judicious spending of taxpayer funds is crucial and both Congress and OMB’s goals are understandable and appropriate. The resulting processes were designed to be deliberate and inflexible, precisely to ensure compliance. Unfortunately, this also means the federal budget and acquisition enterprise is laborious and slow, which is challenging if your goal is to transition and integrate into the IC dynamic and quickly-changing emerging technology like AI. In 2016, the House Committee on the Budget noted “[t]he [budget] process has become so cumbersome, frustrating and ineffective. . .” that it has weakened the power of the purse and Congress’ capacity to govern.⁵⁰ These processes prioritize strict compliance with rules and regulations and predictable cost and schedule⁵¹ over performance and mission effectiveness – a problematic approach if your goal is a high-performing and effective national security community.⁵²

A number of other issues render current budget and acquisition processes inhospitable to AI. First, there is a fundamental mismatch between the timescale over which these technologies evolve and the timescale over which the government can make decisions about how to exploit or respond to that evolution. The executive branch’s three-year budget cycle calls for defined requirements at the beginning of the process⁵³ and changing them at any point along the way is onerous and necessarily implicates cost and schedule.⁵⁴ Similarly, the acquisition process, focused on the purchase of major systems like airplanes and tanks that have clear requirements and milestones, and physical deliverables at the end of the production line, employs Statements of Work (SOWs) to request bids from companies early on, and those SOWs outline detailed requirements. The government expends significant time creating these requirements, and companies are evaluated based on strict compliance with them three or more years later during the execution of the contract. These processes are ill-prepared for software, where major changes can occur over the span of just a few months. AI is fundamentally a software-driven enterprise without a defined deliverable at the end; progress is much more iterative as testing and refinement improve the output. The evolving nature of AI technology is inconsistent with unyielding, overly specific requirements on the front end that are still relevant and useful years later when it is delivered.⁵⁵

50 *Proposed Rewrite of The Congressional Budget Process Summary of Selected Provisions: Hearing Before the H. Comm. On the Budget*, 114th Cong. (2016) (Tom Price, Chairman).

51 Richard Danzig, *Driving in the Dark: Ten Propositions about Prediction and National Security*, Center for New American Security, October 2011 (the US military relies on prediction to forecast needs and influence the design of major equipment and experience shows long-term predictions are consistently mistaken, so acquisition programs should reflect the likelihood of predictive failure and “prepare to be unprepared.”).

52 Allen Schick, one of the foremost experts on the federal budget, has said that to achieve a performing state, governments need a performing budget. He defines a performing state as one that “continuously reads its environment and adjusts how and what it does in response to new information” and notes that there are several tensions that make this difficult, including the pressure to complete budgets on schedule, and the fact that costs of inputs take precedence over results. Allen Schick, *The Performing State: Reflection on an Idea Whose Time has Come but Whose Implementation Has Not*, Inter-American Dev. Bank 1, 4 (Dec. 2002), <https://publications.iadb.org/publications/english/document/The-Performing-State-Reflection-on-an-Idea-Whose-Time-Has-Come-but-Whose-Implementation-Has-Not.pdf>.

53 Patt & Greenwalt, *supra* note 14, at 42.

54 *Id.*, at 49 (“Today, adaptation is essentially forbidden by the necessity of changing program baselines, schedules, program plans, and the associated reprogramming of funds.”).

55 *Rightly Scaled*, *supra* note 35 (“Combined with a rigid adherence to the annual budgeting and appropriations cycle, Congressional processes stymie

In addition, the IC's standard contract vehicles are laborious, slow and rigid. These characteristics do provide some benefits, to include clarity, predictability, and accountability. However, as established, they can become roadblocks when it comes to AI, not only in terms of the dynamic nature of the technology but also, in this case, in discouraging small businesses from engaging with the IC. Small businesses have neither the time nor the resources to wait years to complete a contract and are effectively priced out of the IC market precisely because they cannot afford to wade through difficult and time-consuming contracting processes. Moreover, government contracts are highly susceptible to bid protests from contractors who do not win the contract, which can add exponential delay to the timeline.⁵⁶ And if requirements shift after a contract is completed, the potential for a protest by those who competed and lost based on the old requirements only increases.⁵⁷

Second, the IC's budget process requires spending to fit into a series of discrete sequential steps, represented by budget categories such as research, development, procurement, or sustainment; funds are not quickly or easily spent across these categories.⁵⁸ However, with AI constantly adapting, developing, and adding new functionality, there often is not a clear delineation between research, development, procurement, and sustainment; rather it is a continuous cycle that flows back and forth across these categories rapidly.⁵⁹ New capabilities may be made available to users in very compressed timelines – weeks to months, rather than in years. This makes it extremely difficult to define the beginning and end of a particular step and the result is often budgeting gymnastics, such that one must predict – years in advance – and spend each category of funding accurately or follow laborious and slow processes to adjust funding in the year of execution as a project evolves. Modern practices for fielding AI have outpaced the IC's decades-old approach to budgeting and acquisition. These artificial limitations on the ability to quickly meet the changing needs of a program must themselves evolve to support the full, iterative lifecycle of AI.⁶⁰

Third, most appropriations for discretionary activities expire at the end of each fiscal year, which means programs must develop detailed and predictable spending plans at the beginning of the fiscal year and follow those plans closely. However, when it comes to AI, it is impossible to predict with certainty the timing of developmental breakthroughs, related new requirements, and funding impacts, including funds not being spent as quickly as expected. The potential loss of those funds at the end of the fiscal year can drive last-minute wasteful spending by the IC⁶¹ and causes confusion,

the agility and flexibility needed to match dollars to the pace of rapid technological evolution.”).

56 Annie Palmer & Amanda Macias, Microsoft challenges NSA cloud contract reportedly awarded to Amazon, CNBC (Aug. 12, 2021), <https://www.cnn.com/2021/08/12/microsoft-challenges-nsa-cloud-contract-reportedly-awarded-to-amazon.html>; Katyanna Quach, *Blue Origin sues NASA for awarding SpaceX \$3bn contract to land next American boots on the Moon*, The Register (Aug. 16, 2021), https://www.theregister.com/2021/08/16/blue_origin_lawsuit/.

57 Contract protests are of less concern with sole-source contracts, especially if using the Alpha contracting process, which allows for significant USG-company engagement before finalizing the SOW. Protests also are of less concern with Other Transactions (see below) because OTs are exempt from the Competition in Contracting Act, although may be competed for other reasons. Andrew F. Clements, *A Study of the Alpha Contracting Process and its Effects on Integrated Product and Process Development (IPPD) within Selected Army Acquisition Programs*, Calhoun Institutional Archive of the Naval Postgraduate School 1, 13-20 (Mar. 2002), https://calhoun.nps.edu/bitstream/handle/10945/6071/02Mar_Clements.pdf?sequence=1&isAllowed.

58 This is often referred to as “color of money” because each category requires a different kind, or “color,” of money. See Patt & Greenwalt, *supra* note 14, at 12.

59 MITRE, *supra* note 26, at 24 (recognizing that “software ‘is uniquely unbounded and flexible...’ and the application of legacy accounting rules aligned to multiple appropriations ... is antiquated and restrictive.”).

60 McQuade, *supra* note 45, at 37-8, S76-S78.

61 Jeffrey B. Liebman & Neale Mahoney, *Do Expiring Budgets Lead to Wasteful Year-End Spending? Evidence from Federal Procurement*, Mossavar-Rahmani Center for Business & Government 1, 18-22 (Sept. 2013), <https://www.hks.harvard.edu/sites/default/files/centers/mrcbg/files/mrcbg.fwp.2013-12.Liebman.budgets.pdf> (stating that IT budgets are more likely to be spent at end of year than non-IT budgets); see Elizabeth Harrington, *Washington Free Beacon: Fed Spending Spree Included Millions on Cars, Scooters, Fidget Spinners*, Open Books (Sept. 27, 2018), <https://www.openthebooks.com/washington-free-beacon-fed-spending-spree-included-millions-on-cars-scooters-fidget-spinners/>; see also Luke Chen, *Ditch Use-or-Lose Budgets in the Department of Defense*, Ctr. New Sec. (Mar. 17, 2021), <https://www.cnas.org/publications/commentary/ditch-use-or-lose-budgets-in-the-department-of-defense>; Moshe Schwartz, Cong. Research Serv., R43566, *Defense Acquisition Reform: Background*,

delay, and uncertainty for the private sector. This inefficiency and disruption does not inspire confidence in, or encourage companies to work with, the IC. While expiring funds have the potential to cause these problems broadly, it is completely foreseeable and avoidable when it comes to AI.

Each of these issues results from well-intentioned attempts to buy down the risk of loss or failure and promote accountability and transparency. They require the IC to know with clarity and certainty the solution it seeks in advance of investment and narrowly limiting the IC's ability to change the plan or hastily implement it. But today's technical solutions are constantly evolving, learning, and improving, which both undermines the IC's ability to prescribe a specific solution and, in fact, is incentive for the IC to allow the solution to evolve with the technology. The lack of flexibility and speed in how the IC manages and spends money and acquires goods and services is a core problem when it comes to fully incorporating AI into the IC's toolkit.⁶²

FLEXIBILITY & SPEED: SOLUTIONS

SOOs v. SOWs

A small but impactful step to deal with the slow and rigid acquisition process is to encourage the use of Statements of Objectives (SOO) and Performance Work Statements (PWS) instead of SOWs, when appropriate. As mentioned, SOWs set forth defined project activities, deliverables, requirements, and timelines, which are used to measure contractor progress and success. SOWs make sense when the government understands with precision exactly what is needed, how it should be achieved, and has the time to develop detailed requirements.

SOOs and PWS', on the other hand, are more appropriate when the strategic outcome and objectives are clear, but the steps to achieve them are less so.⁶³ They describe "what" without dictating "how," and therefore encourage and empower industry to propose innovative solutions. In addition, they create clarity about what is important to the government, leading companies to focus less on aggressively low (and often unrealistic) pricing of specific requirements and more on meeting the ultimate outcomes needed in creative ways that align with a company's strengths.

SOOs are useful to streamline the proposal process and add a measure of flexibility, but because they do not bind the government, they must be accompanied by either a PWS or an SOW. Companies often default to SOWs to lock in detailed, tactical requirements. However, in the case of AI, it is preferable to pair an SOO with a PWS, which focuses on strategic requirements. PWS' include more direction than SOOs but do not require extremely specific detail like SOWs.⁶⁴ PWS' are more complicated to write and monitor than SOWs, so they require specially trained and skilled government contracting officers to successfully implement, but they also provide more flexibility to both the government and the company in terms of negotiating the requirements and the potential avenues for achieving them.

Analysis, and Issues for Congress 1, 15 (May 23, 2014) ("the threat that funding will be taken away or that future budgets can be reduced unless funds are obligated on schedule is a strong and perverse motivator. We risk creating incentives to enter into quick but poor business deals or to expend funds primarily to avoid reductions in future budget years.").

⁶² Final Report, *supra* note 7, at 404.

⁶³ Natalie Komitsky, *Will a SOW, PWS, or SOO Work Best to Achieve Your Objective?*, Management Concepts (Sept. 25, 2019), <https://blogs.managementconcepts.com/2019/09/25/will-a-sow-pws-or-soo-work-best-to-achieve-your-objective/>.

⁶⁴ Developing a PWS Handbook, Ctr. Army Lessons Learned, 1, 3-4, <https://www.acq.osd.mil/dpap/ccap/cc/jcchb/Files/Topical/SOWs/09-48.pdf>.

While moving to greater use of the SOO/PWS model is a good start, even that is not as quick or flexible as is desirable for AI. If requirements shift after a PWS is completed, costs may balloon, timelines may stretch, and the contract may be subject to protest by those who competed and lost based on the old requirements.⁶⁵

So several fundamental questions remain. At what point does it make sense to lock in requirements for AI, and at what level? If not before the contract is signed, during the design phase? After a prototype is built? How does this kind of uncertainty impact a legally binding and enforceable contract? Does the government need a more flexible contract model?⁶⁶ Given the global technology race, when it comes to contracting and competition, it is time for the USG to radically re-think “requirements” and create adaptive contracts that allow for evolving needs as circumstances change.

To that end, the SWAP study resulted in DOD creating an adaptive acquisition framework (AAF),⁶⁷ which allows DOD to choose different acquisition approaches for different types of purchases, some more agile and speedy than others. DOD’s rapid Software Acquisition⁶⁸ path allows for continuous feedback loops over the course of a year to achieve a minimum viable product and employs a Capability Needs Statement⁶⁹ to capture high-level mission deficiencies or enhancements, interoperability needs, and other attributes that provide enough information to propose a solution but not so much information as to prescribe an answer. In addition, MITRE recently recommended DOD create an adaptive requirements framework to parallel the AAF.⁷⁰ This framework builds requirements at a portfolio level, as opposed to a tactical level, allowing different capabilities within the portfolio to evolve iteratively. The IC, which has been considering similar approaches, must more quickly learn from and adapt DOD’s experience with the AAF Software Acquisition path and the adaptive requirements framework, both of which hold significant promise for rapid and agile acquisition of AI.⁷¹

In the meantime, the IC must introduce as much flexibility and engagement with a company as possible before a contract is put in place. Using an SOO and PWS together will ensure creative responses by companies to better meet program objectives and is a reasonable first step that will help drive a mindset change around AI.

65 Clements, *supra* note 57, at 13-20; Contract protests are of less concern with sole-source contracts, especially if using the Alpha contracting process, which allows for significant USG-company engagement before finalizing the SOW. Protests also are of less concern with OTs because OTs are exempt from the Competition in Contracting Act, although may be competed for other reasons.

66 See McQuade, *supra* note 45, at S123-26.

67 DAU, Adaptive Acquisition Framework Pathway, <https://aaf.dau.edu/aaf/aaf-pathways/>.

68 DAU, Software Acquisition, <https://aaf.dau.edu/aaf/software/>.

69 DAU, Capability Needs Statement (CNS), <https://aaf.dau.edu/aaf/software/cns/>.

70 See generally MITRE, *Modernizing DoD Requirements: Enabling Speed, Agility, and Innovation*, <https://www.mitre.org/publications/technical-papers/modernizing-dod-requirements-enabling-speed-agility-and-innovation>; See also MITRE, *supra* note 26, at 9 (“The new administration should focus on developing adaptive requirements and budget systems to align with the acquisition pathways.”).

71 The DOD Joint AI Center (JAIC) is also creating a prototype business process to streamline rapid procurement and agile delivery of AI capabilities. This activity, called Tradewind, is a partnership with the Indiana Innovation Institute and will likely yield many helpful lessons for the IC. See JAIC Public Affairs, *JAIC lays the foundation for the Tradewind Initiative in Partnership with Indiana Innovation Institute (IN3)*, JAIC (Feb. 1, 2021), https://www.ai.mil/news_02_01_21-jaic_tradewind_initiative_partnership_in3.html; Tradewind, <https://tradewindfaq.org>.

Additional Authorities

Two additional authorities would help the IC speed and scale its use of AI capabilities immediately: Other Transaction Authority (OTA)⁷² and Commercial Solutions Opening (CSO) authority.⁷³ Both have been in use by DOD and other agencies for years and have proven helpful tools in the government's quest to be more nimble, efficient, and accessible.

OTA allows DOD to more quickly and easily acquire or develop research and prototypes with non-traditional contractors. Other transactions (OTs) are not standard procurement contracts, grants, or cooperative agreements, and therefore generally are not subject to the federal laws and regulations that apply to government procurement contracts (the Federal Acquisition Regulations, or FAR⁷⁴). OTs provide significantly more speed, flexibility, and accessibility than traditional contracts for research initiatives and prototype activities and, therefore, enhance DOD's ability to take advantage of cutting-edge technology like AI.⁷⁵ While they are not a silver bullet, OTs have been used to good effect since 1990 by the Defense Advanced Research Projects Activity (DARPA),⁷⁶ DOD's over-the-horizon research and development organization. However, the IC does not yet have OTA.⁷⁷ The DOD intelligence elements⁷⁸ can use DOD's OTA, but the other IC elements do not have access to it. The IC should be granted OTA as soon as possible, as requested by both the Trump and Biden Administrations.

In addition, DOD, DHS, and GSA have a five-year pilot authority, CSO, which is a simplified and relatively quick solicitation method to award firm fixed price contracts up to \$100 million. CSOs are used to acquire innovative commercial items, technologies, or services that close capability gaps or provide technological advances, among other things, and can be used for anything from research and development to off-the-shelf commercial products. CSOs are an open call for proposals that provide offerors the opportunity to respond with technical solutions of their own choosing to a broadly defined area of government interest. CSOs are considered competitively awarded regardless of how many offerors respond and therefore meet contract competition requirements. The CSO pilot will expire at the end of FY22, so rather than extending the pilot to the IC, CSO authority should be provided permanently to the IC.

Unclassified Sandbox

Speed is of the essence when it comes to the IC's ability to provide insights that inform time-sensitive policy decisions. The predictive nature of the IC's work and the need to forecast outcomes means the IC must be able to acquire AI at the point of need, aligned to the threat. Waiting several years to acquire

72 10 U.S.C. § 2371(b) (2017); See Moshe Schwartz & Heidi M. Peters, Cong. Research Serv., R45521, *Department of Defense Use of Other Transaction Authority: Background, Analysis, and Issues for Congress*, 1, 19-20 (Feb. 22, 2019).

73 DAU, *Defense Commercial Solutions Opening (CSO) Pilot Program*, <https://aaf.dau.edu/aaf/contracting-cone/defense-cso/>, (Section 879 of the FY17 NDAA (PL 114-328, 10 USC 2302 Notes) and Section 880 (PL 114-326, PDF page 315)).

74 The FAR, as will be explained further below, is incredibly complex and laborious, difficult for even acquisition experts to navigate. In many cases, avoiding use of the FAR is the best way to speed an acquisition.

75 OTs are complex and bring other rules and obligations, but OTs exempt the government from certain traditional contract requirements and allow agreements to be structured in numerous ways, including joint ventures, partnerships, and consortia. OTs provide a mechanism for entities to pool resources to facilitate the development of, and obtain, cutting-edge technology, and lower the costs and timelines for engagement with the USG, thereby attracting non-traditional contractors with promising technology to work with DOD. While OTs do have some attendant risks, including diminished oversight and transparency, new and adaptive oversight requirements could substantially mitigate those risks.

76 Defense Advanced Research Projects Agency, <https://www.darpa.mil>.

77 While some in the IC do have the ability to use "Section 8" authority (Sec. 8 of the CIA Act of 1949, 50 USC 3510), which allows, among other things, an agency to contract without using the FAR, Section 8 is very judiciously used given its focus on CIA operations and the original intent of that law being narrowly construed.

78 Transformation & Innovation – Who We Are, *supra* note 28.

these capabilities undermines the IC's ability to fulfill its purpose. But with speed comes added risk that new capabilities will not perform as expected or may completely fail. To address this, the IC should create an isolated unclassified sandbox environment, not connected to operational systems, in which potential IC customers could test and evaluate new capabilities alongside developers in weeks-to-months, rather than years.

To facilitate, Congress should provide the IC with the ability to purchase software up to \$100,000 quickly and easily on a credit card for test and evaluation (T&E) purposes only. In doing so, Congress and the IC would buy down the risk that a rapid acquisition would result in total failure because the sandbox T&E process would provide valuable insight into the tools to quickly determine the potential for success. In addition, the IC would have the ability to test products, consider adjustments, and engage with developers early in the process to speed the ultimate acquisition and increase the likelihood of success.

Single Appropriation

DOD has a pilot program that creates a new single appropriation category for software delivery.⁷⁹ This pilot funds software as a single budget item – allowing the same money to be used for research, production, operations, and sustainment – without regard to those categories. The DoD pilot does not speak to AI specifically and does not extend to the IC but is an important proof of concept.

Like software generally, AI (which consists largely of software) evolves rapidly and requires a unique development approach that does not fit neatly into specific budget categories like research, development, and operations. AI is most effective in practice when tools can be developed side-by-side with operators, evolving and improving based on real-time use and feedback.⁸⁰ Forcing funding to be split between arbitrary budget categories slows the development of those tools and requires more bureaucratic maneuvering than is necessary. This, in turn, makes those tools less timely and useful to the operators who ultimately need them. For these reasons, DOD's single appropriation pilot program is equally applicable to AI and should specifically include it.

Moreover, the same reasons for creating this pilot program in DOD exist in the IC.⁸¹ The pilot's three primary themes include ensuring speed and cycle time, supporting digital talent, and improving on software's unique development cycle. The IC encounters these same issues in working to acquire and develop software in support of national security outcomes and the IC also must take better advantage of software as it pursues its intelligence goals. Despite much of the IC also being part of DOD, IC-specific activities do not fall within this pilot. Extending DOD's pilot to the IC would not only address the primary issues, but it would also allow better interoperability and compatibility between the IC and DOD on joint projects.

79 See MITRE, *supra* note 26, at 24 (This single appropriation category is called BA8); See also Sarah Sybert, DoD Launches Pilot Programs for Software Procurement; Ellen Lord Quoted, ExecutiveBiz, <https://blog.executivebiz.com/2020/08/dod-launches-pilot-programs-for-software-procurement-ellen-lord-quoted/> (Aug. 14, 2020, 10:31 AM).

80 Interim Report, *supra* note 15, at Appendix 1.

81 See U.S. Gov't Accountability Office, GAO-16-464SP, *Principles of Federal Appropriations Law*, 1, 2-3 (2016) ("The very point of a lump-sum appropriation is to give an agency the capacity to adapt to changing circumstances and meet its statutory responsibilities in what it sees as the most effective or desirable way.") (citing *Lincoln v. Vigil*, 508 U.S. 182, 191 (1993)); See also *Hein v. Freedom from Religion Found., Inc.*, 551 U.S. 587 (2007).

No-year Funds

Congress should not use the annual expiration of funds as a lever to control AI development and oversight; instead Congress should consider providing no-year funds for AI development. Congress already routinely provides no-year funding when it makes sense to do so.⁸² The IC should engage with Congress to justify no-year funding for AI due to the uncertain nature of the development of the technology and the difficulty in predicting timelines for completion. Funds that do not expire would allow AI to evolve without arbitrary deadlines, would drive more thoughtful spending throughout the lifecycle of the project, would provide certainty for small businesses, and would eliminate the additional overhead required to manage the expiration of funds annually. In addition, no-year funding would increase Congress' time and capacity to perform oversight of AI activities.⁸³ Given the longer-term nature of this proposal, however, the executive branch also must seek shorter-term solutions.

A less-preferable alternative is to seek two-year funding for AI. Congress has a long history of proposing biennial budgeting for all government activities.⁸⁴ To be effective, biennial budgeting must include both two-year authorizations (approval to spend) and two-year appropriations (availability of funds) – if the approval is for two years but funding remains only available for one year, biennial budgeting is ineffective and will fail, as happened with DOD in the 1980s.⁸⁵ Biennial budgeting proposals continue to go forward regularly in recognition of the significant resources federal agencies would save, as well as the additional time Congress would have for more focused oversight.⁸⁶ Indeed, in 2016, both the House and Senate Committees on Budget reviewed the Congressional budget processes and recommended a series of reform, including biennial appropriations.⁸⁷ These recommendations were made without regard to the type or purpose of funding, but rather were proposed for all appropriations.

Even without a biennial budget, Congress has already provided nearly a quarter of the federal budget with two-year funding. While two-year funding is not a perfect answer in the context of AI, it would at a minimum discourage parties from rushing to outcomes or artificially burning through money at the end of the first fiscal year, and would provide additional time to fulfill the contract.⁸⁸ This is presumably why DOD recently created a new budget activity under their Research, Development, Test and Evaluation (RDT&E) category, which is typically available for two years, for “software and digital technology pilot programs.”⁸⁹ If Congress does not provide no-year funds for AI, it should at the very least include AI in the RDT&E two-year funding category.

82 U.S. Gov't Accountability Office, *supra* note 81, at 2-9 (No-year funding allows money to be spent until expended without a time limit. Multi-year funding allows money to be spent over the course of several years with a defined limit).

83 U.S. Gov't Accountability Office, GAO-/T-AIMD-00-121, *Budget Process: Biennial Budgeting for the Federal Government*, 1, 12-15 (Mar. 10, 2000).

84 *Biennial Budgeting*, Cong. Budget Office 1, 4-12 (Nov. 1987), <https://www.cbo.gov/sites/default/files/cbofiles/ftpdocs/84xx/doc8482/87-cbo-002.pdf>; See U.S. Gov't Accountability Office, *supra* note 82.

85 Robert F. Hale, *Financing the Fight, A History and Assessment of Department of Defense Budget Formulation Processes*, Brookings, <https://www.brookings.edu/research/financing-the-fight-a-history-and-assessment-of-department-of-defense-budget-formulation-processes/>, at 18, 28 (Hale recognizes the “use it or lose it” problem and argues that Congress could extend the life of operating appropriations to allow managers to use year-end spending more responsibly); See also MITRE, *supra* note 26, at 19 (noting that this problem continues today between DOD authorizers and appropriators).

86 MITRE, *supra* note 26, at 21 (“The adoption of a biennial budget cycle would provide the space and time for both DOD and congressional leaders to engage in more detailed information sharing.”).

87 See e.g., Options to Fix the Broken Federal Budget Process: Hearing Before S. Comm. on the Budget, 114th Cong. 3 (2016) (George Everly, Chief Council); Proposed Rewrite of The Congressional Budget Process Summary of Selected Provisions: Hearing Before the H. Comm. on the Budget, 114th Cong. (2016) (Tom Price, Chairman); Recommendations: Hearing before the H. Select Comm. on the Modernization of Congress, (2019) (statement of Chair Derek Kilmer & Vice Chair William Timmons).

88 Jason Miller, Air Force's Next Hack of the Federal Procurement System: One-Year Funding, **Federal News Network** (Dec. 8, 2020), <https://federalnewsnetwork.com/reporters-notebook-jason-miller/2020/12/air-forces-next-hack-of-the-federal-procurement-system-one-year-funding/>. (“eliminating single-year funding is one of the biggest changes that is needed”).

89 John F. Sargent Jr., Cong. Research Serv., R44711, Department of Defense RDT&E: Appropriations Structure 1, 5 (Oct. 7, 2020).

AI Technology Fund

Congress should establish an IC AI Technology Fund (AITF) to provide kick-starter funding for priority community AI efforts, to support the urgent operational need across the IC, and to enable significantly more flexibility to get those projects off the ground.⁹⁰ To be successful, the AITF must have no-year funds, appropriated as a single appropriation, without limits on usage throughout the acquisition lifecycle. In addition, the DNI must exempt the AITF from community taxes during budget reduction activities and align the AITF's activities with annual budget and planning timelines.

The DNI's Director of Science and Technology should centrally manage and allocate the AITF funds, and an AITF Advisory Board of members drawn from the IC elements should oversee the AITF's activities. The AITF Board's responsibilities would include annually helping prioritize a small number of high-potential funding opportunities across the IC based on viable business cases that demonstrate operational need, improved performance, appropriate costs, and manageable risks. Congress could require annual reports on AITF funding priorities, successes and failures, and resource allocation to enable oversight and inform ongoing funding decisions.

The AITF's flexibility and simplicity would incentivize increased engagement by small businesses, better allowing the IC to tap into the diversity of the marketplace, and would support and speed the delivery of AI capabilities to IC mission users. AITF funding for projects would continue for a reasonable period of time, after which IC elements would pick up the funding directly.

ICWERX

As noted, DOD has started tackling these acquisition challenges in many ways, including with new, more nimble innovation organizations.⁹¹ AFWERX is an Air Force innovation organization that drives agile public-private sector collaboration to more quickly leverage and develop cutting-edge technology for the Air Force. AFWERX does this through aggressive use of innovative, flexible, and speedy procurement mechanisms like OTA and the Small Business Innovation Research and Small Business Technology Transfer programs (SBIR/STTR),⁹² which also encourage engagement from small businesses. AFWERX is staffed by acquisition and market research experts who are comfortable using those authorities and understand the market, all critical ingredients of the AFWERX model.

AFWERX has been so successful that it has sparked other WERX projects – some with similar approaches, like SPACEWERX, and some that are not government activities but are rather run as non-profits, such as DEFENSEWERX and SOFWERX. And while the IC's needs are not identical, the IC should look to the AFWERX model in terms of aggressively using acquisition flexibilities to speed contracting – lowering the barrier to entry for small businesses to tap into the diversity of the marketplace more effectively – and building a cadre of officers with deep acquisition and market research expertise for use by the entire IC.⁹³

⁹⁰ Nat'l Sec. Comm'n Artificial Intelligence, *supra* note 7, at 673-75 (The ETF should incorporate features of both the Artificial Intelligence Critical Applications Fund for the IC, proposed by the 2021 NSCAI Final Report and the government's Technology Modernization Fund); See H.R. 2227, 115th Congress (1st Sess. 2017).

⁹¹ There has been some criticism that DOD has too many innovation organizations and they should be consolidated. See Melissa Flagg & Jack Corrigan, *Ending Innovation Tourism*, CSET (July 2021), <https://cset.georgetown.edu/publication/ending-innovation-tourism/>. Even if true, the WERX construct has been successful and should be maintained; See MITRE, *supra* note 26 at 38-9.

⁹² *About, SBIR STTR*, <https://www.sbir.gov/about>.

⁹³ Ctr. Strategic & Int'l Stud., *supra* note 38, at 28.

Through the AIM initiative, the IC already has the foundation to create its own WERX-like organization – ICWERX – modified to meet the IC’s needs. The AIM initiative created an AI innovation hub, called the iHub, which connects AI and machine learning experts and activities across the IC. It serves as a coordinating function for cross-IC pilots and projects, provides an entry point for engagement with the private sector, and is intended to become an innovation incubator and accelerator for the IC.⁹⁴ The AIM initiative has tight connections to DOD counterparts, so iHub personnel could quickly partner with AFWERX specialists to learn about the WERX model and adapt it for IC purposes. The iHub should broaden its focus to take on a WERX mission and portfolio, building IC acquisition and market research expertise, and creating best practices for using flexible authorities and engaging with the private sector.⁹⁵

In addition to housing technology and acquisition specialists, ICWERX must employ an integrated, multi-disciplinary IC team that includes legal, financial, operational, and innovation experts, each of whom play a critical role in the IC’s adoption of AI. ICWERX should leverage its connections to other IC innovation labs to coordinate, collaborate, and deconflict activities. To facilitate the IC’s engagement with the private sector, ICWERX should provide initial funding against its projects through the AITF, described above. Finally, ICWERX should house the IC’s unclassified sandbox to encourage IC customers and private sector companies to buy down the risk by testing and evaluating capabilities before entering into more binding contractual arrangements. The National Geospatial-Intelligence Agency recently opened Moonshot Labs, an unclassified innovation center that will foster collaboration among government, industry, and academia. Moonshot Labs will have best practices and lessons learned that should inform ICWERX and the sandbox proposal.⁹⁶

ICWERX would serve as a sorely-needed accessible and responsive “front door” for prospective partners and vendors who are unfamiliar with the IC, encourage collaboration and unity of effort across IC AI activities, and add a vital tool to the IC’s toolkit to more quickly leverage and scale cutting-edge AI. ICWERX would complement the work of IARPA, DARPA,⁹⁷ In-Q-Tel, and others. As discussed previously, while these organizations are known for driving innovation and cutting-edge solutions in government, their activities are not intended to facilitate broad IC collaboration with the private sector to quickly transition existing AI tools at scale.⁹⁸ To the extent the IC seeks to accomplish something within IARPA, DARPA, or In-Q-Tel’s scope of work, the IC should leverage those organizations.

94 Ctr. Strategic & Int’l Stud., *supra* note 38, at 19 (the iHub could also play the role envisioned by CSIS’ Task Force on Maintaining the Intelligence Edge, calling for an IC SkunkWorks).

95 This would also help with proper and consistent administration of OTs and provide better oversight of IC use of OTs. Schwartz & Peters, *supra* note 72.

96 NGA celebrates opening of Moonshot Labs in St. Louis July 23, National Geospatial Intelligence Agency (July, 19, 2021), https://www.nga.mil/news/NGA_celebrates_opening_of_Moonshot_Labs_in_St_Loui.html.

97 DARPA, Defense Advanced Research Projects Agency, <https://www.darpa.mil>.

98 Some have questioned whether DARPA, IARPA, and In-Q-Tel’s authorities might be extended to the rest of the IC. IARPA and DARPA largely rely on research and development authorities – neither available nor appropriate for the broader IC. In addition, DARPA highly agile (more so than IARPA) due to its autonomy over personnel, security, and contracting – DARPA controls how and how fast its activities proceed, and both DOD and Congress allow DARPA that latitude. In addition, there is an R&D ethos that affixes to DARPA and IARPA, that insulates them from a significant amount of the daily oversight and tasking that is common for the rest of the IC. Separately, In-Q-Tel is an independent organization not constrained by government process or procedure and can use appropriated funds in multi-year increments – in other words, their funds do not expire each year, as recommended above.

ACTIONS: FLEXIBILITY & SPEED

1. The ODNI should require certain types of programs to use the SOO/PWS model by default.
 - a. Identify criteria for programs that should use an SOO/PWS rather than an SOW.
 - b. Identify SOO/PWS experts in each organization who can mentor less experienced officers and identify appropriate additional training.
 - c. Accelerate creation of an AAF that includes a software acquisition path, learning from DOD's experiences.
2. Congress should immediately provide OT and CSO authority to the IC.
3. The ODNI should create an unclassified, stand-alone IC sandbox housed in ICWERX where IC customers can test and evaluate AI capabilities quickly to buy down risk and speed up acquisition.
 - a. Congress should authorize the IC to quickly and easily spend for T&E purposes up to \$100,000 per capability to facilitate the sandbox evaluation process.
4. Congress should extend DOD's single appropriation for software pilot to the IC.
5. The ODNI should work with Congress to create no-year funds for AI.
 - a. At a minimum, all AI must be designated as part of RDT&E 2-year funds.
6. The ODNI should work with Congress to establish an IC AI Technology Fund, administered by the DNI's Director of Science & Technology.
 - a. Must have appropriated funds that are no-year, no-color, and non-taxable.
 - b. Define what qualifies as "AI" for purposes of this fund.⁹⁹
7. The ODNI should use the AIM iHub as the foundation for an ICWERX, modeled after AFWERX, to increase speed, agility, and ease of engagement with private sector innovators and small businesses developing AI tools and technologies.
 - a. Create a loose partnership with AFWERX, to learn from their experts and leverage their strengths while building the expertise of the IC's acquisition professionals regarding the use of authorities and market research.
 - i. Consider co-location of small ICWERX within AFWERX to start.
 - ii. Prioritize and hire additional acquisition and market research experts.
 - b. Leverage and increase iHub/ICWERX connections to other IC innovation labs and centers to share expertise and consult. This will enable a better understanding of the IC's capabilities and needs, better training of the IC's workforce, better use of IC authorities, better partnering with the private sector, and ultimately faster time to mission for emerging technologies.
 - c. Use AI Technology Fund to fund specific activities within ICWERX.
 - d. House unclassified IC sandbox in ICWERX for quick and easy T&E of new AI capabilities.

⁹⁹ The DNI's AIM strategy, which defines AI, should be the starting point. *See generally The Aim Initiative*, Director of National Intelligence, <https://www.dni.gov/files/ODNI/documents/AIM-Strategy.pdf>.

COMPLEXITY: ISSUES

Budget and acquisition laws and processes are each formidable in their own right, and when taken together they introduce even greater complexity and uncertainty in interpretation and application.¹⁰⁰ The rules are a labyrinth of requirements, regulations, and processes meant to ensure appropriate spending and oversight of taxpayer dollars,¹⁰¹ but even long-time acquisition professionals have trouble understanding all of the rules and flexibilities, how and when to use them, and the associated risks.

As one example, the FAR is not commonly considered flexible, but experts will tell you there is quite a bit of flexibility in it.¹⁰² In fact, the government has created a tool called the TechFAR, which highlights flexibilities in the FAR and can help agencies better navigate the acquisition regulations for digital service projects.¹⁰³ Unfortunately, the FAR is so complex and difficult to understand¹⁰⁴ that many acquisition professionals remain unaware of the FAR's flexibilities or unwilling to use them in the face of uncertainty.¹⁰⁵

In addition, actions proposed in this report, such as the use of SOOs and PWS', require highly experienced officers who can appropriately convey government expectations and set reasonable metrics and timelines even when specific steps and details remain unknown. As more flexibility and uncertainty is introduced, it becomes increasingly critical for acquisition professionals to be knowledgeable and comfortable with the tools and levers they must use to appropriately manage and oversee contracts.

More generally, there is no quick or easy way for practitioners to keep up with changes in acquisition rules, and the IC has a distributed approach that allows each IC element to use its various authorities independently, rather than cohesively and harmoniously. So the complexity extends across agency lines and there is ample opportunity for confusion.

Many departments and agencies have acquisition courses and materials aimed at demystifying the acquisition authorities. For example, DOD has the Defense Acquisition University, which focuses on defense acquisition strategies.¹⁰⁶ The Defense Intelligence Agency (DIA) has partnered with the University of Maryland's Applied Research Laboratory for Intelligence and Security to create a new emerging technology acquisition course.¹⁰⁷ The Department of Homeland Security (DHS) has a Procurement Innovation Lab that includes a Boot Camp focused on innovative acquisition techniques.¹⁰⁸ And the Federal Acquisition Institute (FAI),¹⁰⁹ which supports career development and strategic human capital management for the acquisition workforce, has an online knowledge management portal for innovative business practices and technologies, called the Periodic Table of Acquisition Innovations,¹¹⁰ to help federal officers navigate acquisition and budget intricacies. However, there is no connection between these programs, no clear path for IC officers to participate, and no reward for doing so.

100 See Hale, *supra* note 14, at 29-30, 34; See also Patt & Greenwalt, *supra* note 14, 25-6 (describing complexity of acquisition regulations).

101 See Patt & Greenwalt, *Id.* at 26 ("Thus, by 1971, the triad of budgeting, requirements, and acquisition processes was finally in place. All revolve around prediction of the future").

102 See MITRE, *supra* note 26, at 13.

103 See e.g., *TechFAR Handbook*, TechFAR Hub, <https://techfarhub.cio.gov/handbook/> (highlighting the flexibilities in the FAR that can help agencies implement a series of "plays" drawn from the Digital Services Playbook); Digital Services Playbook, U.S. Digital Service, <https://playbook.cio.gov/>.

104 Noah Shachtman, *Pentagon's Craziest PowerPoint Slide Revealed*, *Wired* (Sept. 13, 2018), <https://www.wired.com/2010/09/revealed-pentagons-craziest-powerpoint-slide-ever/>.

105 McQuade, *supra* note 45, at 31.

106 *Defense Acquisition University*, DAU, <https://www.dau.edu>.

107 *TLA 2021 Online Course: Acquisition of Emerging Technologies*, Applied Res. Lab. Intelligence & Sec., https://www.arlis.umd.edu/tla2021_aet.

108 *Procurement Innovation Lab*, Dep't Homeland Sec. (Mar. 5, 2021), <https://www.dhs.gov/pil>.

109 *About FAI*, Fed. Acquisition Inst., <https://www.fai.gov/about/about-fai>.

110 *Periodic Table of Acquisition Innovations*, Fed. Acquisition Inst., <https://www.fai.gov/periodic-table>.

Even when acquisition officers firmly grasp the authorities they have, there is little incentive to wade through this complexity and push the boundaries on flexibility. Acquisition officers are rarely brought into substantive mission discussions early and are often not made to feel part of the team. They are viewed as a means to an end and engaged only when a contract must be completed. This results in acquisition officers having a limited understanding of the goals of a contract and minimal interest in the outcome. As if to reinforce this, they are evaluated and rewarded for obligating funds, not for the success of a contract's outcomes.

Moreover, officers who bind the U.S. government with their signature are held personally liable and professionally responsible for their contracting actions, creating a risk-averse and compliance-focused culture with little incentive for innovation or creativity in contracting approaches. Using traditional budget and acquisition processes increases confidence in the ability to timely spend expiring funds in accordance with the rules, so acquisition officers understandably prefer to use those well-understood, less risky paths, and more flexible authorities are not used to the fullest extent possible.

COMPLEXITY: SOLUTIONS

De-mystify Current Authorities

The IC must understand and fully use the authorities it already has. While there is much complaining about a lack of flexible authorities in the IC (and a real need for legal reform), there is flexibility in existing rules that has not been fully utilized.¹¹¹ Unfortunately, the IC has not prioritized the development or hiring of people with government acquisition and contracts expertise, so there are insufficient officers who know how to use the existing authorities and those who do are overworked and underappreciated. The IC must redouble its efforts to increase its expertise in and support the use of these flexibilities in a variety of ways, to include with the required cultural changes described above.¹¹²

DEVELOP PARTNERSHIPS

The IC must create formal partnerships and consult with existing US government experts¹¹³ in places like the General Services Administration's Technology Transformation Services (TTS) and FEDSIM.¹¹⁴ TTS is focused on helping the US government take better advantage of modern technology and houses a center of excellence focused specifically on accelerating the government's adoption of AI.¹¹⁵ FEDSIM provides assisted acquisition services and helps build innovative acquisition solutions for federal agencies. The ODNI and the IC must formally partner with TTS and FEDSIM to identify and incorporate their best practices and lessons learned and to leverage their expertise in acquiring and adopting AI.

DOD's Joint AI Center (JAIC) also has built significant acquisition expertise the IC must leverage.¹¹⁶ In addition to its increased focus on educating the DOD workforce on transformative and agile

¹¹¹ See MITRE, *supra* note 26, at 13.

¹¹² *Id.* at 10.

¹¹³ See also MITRE, *supra* 26, at 12 (DOD must build greater cross-governmental partnerships ... creating opportunities to learn and share our knowledge).

¹¹⁴ *Technology Transformation Services*, Gen. Serv.'s Admin. (Feb. 23, 2021), <https://www.gsa.gov/about-us/organization/federal-acquisition-service/technology-transformation-services>; Our Work, FEDSIM, <https://fedsim.gsa.gov/ourwork/>.

¹¹⁵ *Artificial Intelligence*, Ctr.'s Excellence, <https://coe.gsa.gov/coe/artificial-intelligence.html>.

¹¹⁶ The JAIC has been successful for many reasons, including that the Deputy Secretary of Defense charged the JAIC with taking risk and made them a direct report, signaling how important this effort is to the department.

acquisition strategies, the JAIC recently launched an OTA project called Tradewind.¹¹⁷ Tradewind seeks to streamline, simplify, and accelerate AI adoption through agile and efficient business processes. Through Tradewind, DOD is working transparently within DOD, with the private sector, and with academia to create a whole-of-nation approach to support DOD's AI innovation needs. The IC must immediately partner with the JAIC more fully to learn from and tap into its expertise and the Tradewind activity so that the IC can similarly streamline, simplify, and accelerate AI adoption.

Finally, the IC should increase joint duty rotations in this area, exchanging acquisition personnel between IC elements for several years at a time, to better integrate, leverage, and impart acquisition expertise across the IC.¹¹⁸

PRIORITIZE TRAINING & EDUCATION

The IC must prioritize training and education. The IC's acquisition professionals must have a strong understanding of the existing acquisition processes, rules, and mechanisms that support the acquisition and adoption of AI.¹¹⁹ And while deep acquisition expertise is not necessary for everyone, it is important for lawyers, operators, technologists, and innovators to have a good understanding of acquisition processes and procedures (and acquisition law for the lawyers), including what it takes to complete a contract and the role they play in helping to reach successful outcomes throughout the process. This training is equally important for supervisors and managers as it is for practitioners to help eliminate the potential for surprise, confusion, and delay throughout the process. Collaboration and understanding across these professions and up and down the chain of command will result in more cohesive, speedy, and effective outcomes.

To that end, the ODNI should work with the many government and IC organizations creating acquisition education programs to develop avenues for all IC officers to attend, either in individual courses or as a program of courses to grow deep expertise. The ODNI should strengthen continuing education requirements for specific IC skill fields and reward those who complete a certain number of courses and maintain proper certification with incentive pay – a higher level of pay for these desirable skills. The National Intelligence University (NIU),¹²⁰ the IC's only accredited, federal degree-granting university, should prioritize the creation of courses related to AI, including acquisition, through focused workshops and the creation of a certificate program. The NIU also could work with the existing acquisition schools and courses to create an education pathway for IC officers.

HIRE EXPERTISE

The IC must prioritize and hire deeply experienced and highly qualified experts in government acquisition.¹²¹ These experts should not only include very senior short-term hires, which is often the case, but also should include officers at more junior levels who can become permanent employees with room to grow and the opportunity for a long career in the IC.¹²² Government acquisition is

117 JAIC Public Affairs, *supra* note 71; Tradewind, *supra* note 71.

118 *Joint Duty – About*, Off. Dir. Nat'l Intel., <https://www.dni.gov/index.php/careers/joint-duty/about-joint-duty>.

119 See MITRE, *supra* note 26, at 29-30 (DOD also needs to update its acquisition education and may provide useful models).

120 *About NIU*, Nat'l Intelligence Univ., <https://ni-u.edu/wp/about-niu/>.

121 The term 'highly qualified experts' is used here both generically and as defined in Intelligence Community Directive 623, Appointment of Highly Qualified Experts, which allows the IC to hire HQEs at more competitive salary rates for a limited period of time. *Intelligence Community Directive Number 623*, Office Dir. Nat'l Intelligence (Oct. 16, 2008), https://www.dni.gov/files/documents/ICD/ICD_623.pdf; See Chapter 6: Technical Talent in Government, Nat'l Sec. Comm'n A.I. (Mar. 2021), <https://reports.nscai.gov/final-report/chapter-6/>.

122 These experts could be hired by individual IC elements or as an IC pool of talent, akin to the U.S. Digital Service, and then be deployed across the IC to assist with individual acquisition activities, and help guide and mentor IC acquisition organizations.

a specialized skillset that takes years to master; cycling experts in and out every few years all but guarantees that the IC will have to constantly train new people without building any bench of in-house expertise. And educating and training the current workforce is not a near-term solution. IC elements must make better use of direct hire authority to quickly bring in government acquisition and contracts expertise from outside government and should leverage the national labs, federally funded research and development centers, and university-affiliated research centers to cultivate talent.¹²³ There is a severe shortage of acquisition expertise in the IC today. Without additional experts to quickly drive the actions in this report, the IC will stay right where it is when it comes to AI.

FINALIZE REFERENCE BOOK

Practitioners should have an easily accessible reference book to more quickly discover relevant authorities, understand how to use them, and find consultants who are community experts. FAI's Periodic table is a good start but extremely high-level and not specific to the IC. The White House Office of Science & Technology Policy and OMB also published an innovative contracting playbook in 2014 that could be helpful but needs an update.¹²⁴

In recent years, as part of the IC's Acquisition Agility Initiative,¹²⁵ the ODNI led the creation of an IC Acquisition Playbook that describes common IC acquisition authorities, practices and usages. The goal was to ensure IC acquisition professionals understood and could fully utilize available options for acquisition and contracting across the IC. The current version of the Playbook is an excellent start; however, it is overly focused on DOD authorities and also provides only a general understanding of the authorities; it does not provide a detailed explanation of creative usage or points of contact for questions. Moreover, the Playbook is not yet easily accessible to IC officers. The ODNI should further develop and disseminate this Playbook as a quick win for the IC.

Incentivize Behavior

To encourage creative and innovative acquisition practices, the IC must develop a team approach that focuses on collaboration between acquisition officers, innovators, operators, technologists and others,¹²⁶ and aligns incentives with desired processes and outcomes. Today, acquisition professionals are often brought into projects only in transactional ways, when contracts must be completed or money must be obligated, for example.¹²⁷ So those officers do not have a full picture of the overarching project, are not aware of the desired outcomes, are not part of developing the solutions, and have no real investment in the project's success. Bringing acquisition professionals into the conversation earlier will allow the entire project team to better understand the acquisition processes, nuances, and options, and will create in acquisition professionals a more vested interest in the success of the contract. This, in turn, will inspire more novel and unconventional acquisition ideas to achieve those outcomes.

In addition, in practice, there is an intense focus on acquisition officers obligating all of their funds by the end of the fiscal year. Spending money in a timely manner is important, but it is not the only

¹²³ *Policy, Data, Oversight*, Off. Pers. Mgmt., <https://www.opm.gov/policy-data-oversight/hiring-information/direct-hire-authority/#url=Governmentwide-Authority>.

¹²⁴ See MITRE, *supra* note 26, at 13.

¹²⁵ Now called Reimagining Acquisition.

¹²⁶ See MITRE, *supra* note 26, at 13 ("A discipline of collaboration aims to produce partnerships across functional domains, reaching out beyond the usual suspects, and making a determined effort to remove barriers to participation").

¹²⁷ Based on interviews with current acquisition professionals.

valuable measure of success for these officers. Consistent with the team approach, evaluation standards should also focus on effective engagement and collaboration with internal and external stakeholders, having a good understanding of the mission need, using decision processes that drive creative consideration of alternatives and options, and successful delivery of mission outcomes. If a project is not successful, officers should not automatically receive a negative review, rather consideration should be given to whether the risk was reasonable and what they learned from the experience that may inform future success. These are more subjective measures than are traditionally used but would critically shift the focus to delivering effective outcomes against a contract, rather than merely obligating money. And this shift would incentivize acquisition professionals to use the expertise of innovation offices like ICWERX to dive more deeply into creative solutions that ultimately may be more effective than traditional options.

Reorienting evaluation standards in this way requires several prerequisites: training for acquisition professionals in these new skills; a requirement for project teams to engage with acquisition professionals early; supportive, knowledgeable, and risk-tolerant leadership; and authority for acquisition professionals to make decisions and changes based on evolving facts and conditions. These are not small requirements and they all demand a culture that embraces change and taking calculated risk.

Lastly, the ODNI must highlight and reward the desired behaviors. The ODNI should reinvigorate the IC acquisition awards to publicly celebrate successful outcomes, and prioritize IC mission team awards that include acquisition professionals, in recognition of their impact on the ultimate success of the mission.

ACTIONS: COMPLEXITY

1. The IC must fully understand and use flexible authorities it already has for both budget and acquisition.
 - a. Partner with and leverage USG-wide expertise (like TTS and FEDSIM) to increase IC's knowledge base.
 - i. Strengthen partnership with the DOD JAIC acquisition team, including engagement on its Tradewind initiative.
 - b. Develop comprehensive IC training and development program for everyone from acquisition, contract and procurement officers to lawyers to operators and others associated with the acquisition. The ODNI should:
 - i. Partner with existing programs at DIA, DHS, DOD and elsewhere to enable IC officers to participate in existing courses.
 - ii. Charge NIU with helping to address the IC education gap in acquisition and emerging technology through workshops, a certificate program, or by linking courses at DIA, DHS, DOD and elsewhere.
 - iii. Create courses for various levels of officers, including managers and supervisors, who should understand acquisition requirements and processes at a high level and the impact that work has on the mission.

- iv. Ensure training explores existing FAR flexibilities, OTA, SBIR/STTR, Statements of Objectives, market research, and data rights.
- v. Ensure contractors working on program management or acquisition activities have sufficient expertise and meet training standards.
- vi. Strengthen continuing education requirements for all IC acquisition professionals to keep skills fresh and reward those who stay current on critical skills with incentive pay.
- c. Prioritize and speed the recruitment and hiring of highly qualified acquisition professionals using existing direct hire authorities.
- d. Further develop and accelerate promulgation of the IC Acquisition Playbook and lessons learned to every IC acquisition professional to both ensure the ability to use authorities and to encourage a community approach to usage.
 - i. Start by posting the current Playbook on the ODNI website so it is easily accessible to IC officers.
 - ii. Link IC Playbook to FAI Periodic Table for additional authorities and examples.
- 2. The ODNI should align incentives with desired behavior.
 - a. Create a collaborative team approach that brings acquisition professionals into projects earlier in the process.
 - b. Evaluate performance of IC acquisition professionals, their supervisors and managers against processes and outcomes, in addition to funding obligation rates.
 - i. Include focus on engagement and collaboration with customers and partners, meaningful decision processes, delivering mission outcomes, and learning from failures.
 - ii. Require project teams to engage with acquisition professionals early.
 - c. Reinvigorate IC-wide acquisition and mission awards that recognize acquisition professionals' mission impact, innovation, and cross-community collaboration.

Between the rigid budget and acquisition processes and confusion about how to apply them, there is very little ability for the IC to take advantage of a fast-moving field that produces new and updated technology daily. It is past time for a paradigm shift away from the IC's traditional, linear processes to more iterative and dynamic approaches that will speed and transform the way the IC purchases, integrates, and manages the use of AI. However, doing so brings added risk, not just to cost, schedule and performance but also to Congress and OMB, both of which risk losing some control and insight into spending, which goes to the core of their power. Therefore, in proposing to increase flexibility for the IC, it is critical to create in parallel new mechanisms to assess and manage risk and to enhance Congress and OMB's oversight of the IC in new and adaptive ways. We address these topics next.

Risk – A Scary Prospect

By its very design, AI tools are evolving, learning, and adapting all the time. This constant evolution can bring substantial improvement and benefit in many cases, but also can introduce uncertainty and mistakes; working with AI, like other emerging technologies, necessarily involves some level of risk and failure.¹²⁸ Ultimately, some technology will not work as expected, some companies will not survive, and some delays and unexpected deviations will occur. For each incredible new invention, there are hundreds of brilliant ideas that have failed. But to entrepreneurs and innovators, “failure” is not a bad word. Indeed, failed ideas are often critical steps in the learning process that ultimately lead to a successful product; without those prior failed attempts, the final product might never be created.¹²⁹ As former President of India A.P.J. Abdul Kalam once said, “FAIL” should really stand for “First Attempt In Learning.”¹³⁰

The USG, however, is not Silicon Valley; it does not consider failure a useful part of the process, especially when it comes to national security activities and taxpayer dollars. There is rarely a distinction made within the government between big failures, which may have a lasting, devastating, and even life-threatening impact, and small failures, which may be mere stumbling blocks with acceptable levels of impact that result in helpful course corrections. The HPSCI STAR report emphasizes this, noting that “program failures are often met with harsh penalties and very public rebukes from Congress, which often fails to appreciate that not all failures are the same.”¹³¹

The dynamic and evolving nature of AI means solutions may change over time, increasing costs and adding delay, and there will very likely be some technology that just does not work as intended. Such changes will not always result in significant loss, but in the IC they often lead to congressional hearings, inspector general reports, performance evaluation downgrades, negative reputational effects, and even personal liability.¹³² Officers seek to avoid these results at all costs, so the entire system across the executive and legislative branches creates a risk-averse culture that is terrified of failure.

While there are certainly some risks the USG should not take – where billions of dollars or lives are at stake, for example – there are also some areas that should allow for greater risk-taking.¹³³ For example, when purchasing technology that will require certain but unknown adjustments, impacting to a lesser degree both cost and timing, to deliver the desired results. No one in the USG wants to incur additional costs or delay, or lose taxpayer dollars. But again as the HSPCI highlights, “[e]specially with cutting-edge research in technologies. . . early failures are a near certainty and, so long as they are not due to negligence, should be considered learning opportunities. In fact, failing fast and adapting quickly is a critical part of innovation.”¹³⁴ There is a vital difference between an innovative project that fails and a failure to innovate. The former teaches us something we did not know before, while the latter is a national security risk. The IC must rethink its tolerance for taking risks in a field where failures are embraced as part of the key to future success. The IC must shift from a culture that punishes even reasonable risk to one that embraces, mitigates, and owns

128 Nat’l Sec. Comm’n Artificial Intelligence, *supra* note 7, at 132.

129 *First Attempt in Learning: Why Failure is Key to Innovation*, Twenty One Toys, <https://twentyonetoy.com/blogs/learning-from-failure/fail-first-attempt-in-learning-why-failure-key-innovation>.

130 A.P.J. Abdul Kalam, Goodreads, <https://www.goodreads.com/quotes/1015959---if-you-fail-never-give-up-because-f-a-i-l-means>.

131 Rightly Scaled, *supra* note 35, at 33.

132 31 USC § 3528 (1998).

133 Nat’l Sec. Comm’n Artificial Intelligence, *supra* note 7, at 133.

134 Rightly Scaled, *supra* note 35, at 33.

it.¹³⁵ This can only be done with a systematic, repeatable, and consistent approach to making risk-conscious decisions.

Today there is no cross-IC mechanism for thinking about risk in this area, let alone for taking it. The IC is typically risk-averse and prefers not to introduce any new risk. That is, of course, not realistic. And it also is not the standard the IC meets today in its operations. The IC is constantly faced with a multitude of operational risks: that its officers, sources, or methods will be exposed, that it will miss (or misinterpret) indications of an attack, or that it will otherwise fail to produce the intelligence policymakers need at the right time and place. Yet in the face of such serious risks, the IC proactively and aggressively pursues its mission. It recognizes that it must find effective ways to understand, mitigate, and make decisions around risk, and therefore it takes action to make sure potential ramifications are clear, appropriate, and accepted before any failure occurs. This recognition must also be applied to the ways in which the IC acquires, develops, and uses new technology.

The IC cannot be paralyzed by a zero-risk tolerance that is neither attainable nor desirable. Indeed, there is a substantial risk inherent in being too constrained by risk-aversion. In the context of AI specifically, the threat of losing the race with China is of significant concern and weighs in favor of taking more AI-related risks.¹³⁶ For these reasons, the IC must recognize and affirmatively work to understand, mitigate, and make decisions that manage risk, while at the same time understanding and accepting reasonable consequences that result from related failures.

When considering new activities or approaches, each IC element manages risk through its own lens and mechanisms, if at all. Some IC elements have created internal risk assessment frameworks to help officers understand the risks of both action and inaction, and to navigate the decisions they are empowered to make depending on the circumstances. These frameworks provide confidence that if an activity goes wrong, supervisors all the way up the chain will provide backing as long as the risk was reasonable, well-considered and understood, and the right leaders approved it. However, these frameworks are independent and speak only to individual IC element risks – they are not IC-wide or even necessarily consistent with each other.

To drive interoperability and integration, the broader IC must have an overarching risk framework that ties together these disparate approaches and provides consistency and understanding across the IC elements. Each IC element can and should tailor the IC framework to its own unique needs and missions as needed but, at a minimum, each IC element should have the benefit of a robust and repeatable framework that can be used across the lifecycle of an AI system if the IC is expected to be innovative, take acceptable risks, and ultimately take full advantage of AI. While risk assessments are often not precise instruments of measurement, reflecting the quality of the data used, the skills and expertise of those conducting the assessments, and the subjective interpretation of the results,¹³⁷ they are a key part of effective risk management and facilitate decision-making at all levels.

The unique nature of the IC does not lend itself to a one-size-fits-all approach: the IC is comprised of 18 different elements, each with similar and overlapping, but not identical, missions, roles, and authorities, and different threats and vulnerabilities. A flexible, common framework for considering

¹³⁵ See MITRE, *supra* note 26, at 28 (noting that failure should be considered a learning process that leads to improvement rather than punishment).

¹³⁶ Final Report, *supra* note 7 at 113 (“The IC needs to balance the technical risks involved in bringing new technologies on line and quickly updating them with the substantial operational risks that result from not keeping pace”).

¹³⁷ *Guide for Conducting Risk Assessments*, National Institute of Standards and Technology (Sept. 2012), <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf>.

risks that IC elements can tailor to their own processes and programs would significantly improve the workforce's ability to understand acceptable risks and tradeoffs, produce comprehensible and comparable risk determinations across the IC, and provide policy-makers the ability to anticipate and mitigate failure and unintended escalation.

ELEMENTS OF AN AI RISK FRAMEWORK

A common IC AI risk framework¹³⁸ should inform and help prioritize decisions regarding AI from acquisition or development, to deployment, to performance in a consistent way across the IC to enable community activities. As a starting point, the IC should create common AI risk management principles that include clear and consistent definitions, thresholds, and standards. These principles should drive a repeatable risk assessment process that each IC element can tailor to its individual needs, and should promote policy, governance, and technological approaches that are aligned to risk management.

The successful implementation of this risk framework requires a multidisciplinary approach, in some cases involving leaders from across the IC, and always involving experts from all relevant functional areas and managers who can ensure vigilance in implementation.¹³⁹ A whole-of-activity methodology that includes technologists, collectors, analysts, innovators, security officers, acquisition officers, lawyers and more, is critical to ensuring a full 360-degree understanding of the opportunities, issues, risks, and potential consequences associated with a particular action, and to enabling the best informed decision.

Given the many players involved, each IC element must strengthen internal processes to manage the potential disconnects that can lead to unintended risks. Internal governance processes should include an interdisciplinary Risk Management Council (RMC) made up of senior leaders from across the organization. The RMC should establish clear and consistent thresholds for when a risk assessment is required, recommended, or not needed given that resource constraints will not allow all of the broad and diverse AI activities within organizations to be assessed. These thresholds should be consistent with the IC risk management principles so that as IC elements work together on projects across the community, officers have similar understandings and expectations.

The risk framework itself should provide a common taxonomy and process to:

- Understand and identify potential failures, including the source, timeline, and range of effects.
- Analyze failures and risks by identifying internal vulnerabilities or predisposing conditions that could increase the likelihood of adverse impact.
- Evaluate the likelihood of failure, taking into consideration risks and vulnerabilities.
- Assess the severity of the potential impact, to include potential harm to organizational operations, assets, individuals, other organizations, or the nation.
- Consider whether the ultimate risk may be sufficiently mitigated or whether it should be transferred, avoided, or accepted.

¹³⁸ See Appendices A-E (while this framework is written through the lens of AI, it could easily be adapted to cover risk more broadly).

¹³⁹ Benjamin Cheatham, *Confronting the risks of artificial intelligence*, McKinsey & Company (Apr. 26, 2019), <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/confronting-the-risks-of-artificial-intelligence>.

As each IC element assesses potential failures, vulnerabilities, likelihood, and severity, it should consider the risks relevant to the individual AI capability or toolset. AI-related risks may include:

- First order risks (starts chain of events)¹⁴⁰:
 - Technology failure: development, testing, or operational performance is not completed as expected or does not work in practice
 - Data & Algorithms: insufficient, incorrect, bias or other faulty data inputs, incorrect storage, access, retention, dissemination, disposition of data
 - Adversarial Attack: to deny, disrupt, degrade or destroy AI capabilities or infiltrate organization
 - Security: infiltration of systems; insider access
 - Supply chain security or integrity breach: required inputs are intentionally or unintentionally damaged or compromised such that technology does not work or is not trusted
 - Operational need: ability to execute mission without technology
 - Partner failures: inability to provide reliable, quality inputs; survival of company
 - Cost increases: leading to insufficient funds or overrun
 - Schedule slips: delays may render technology less useful or obsolete
 - Policy: unclear or insufficient policy and regulation of technology and usage
- Second order risks (the result of first order risks)¹⁴¹:
 - Workforce: personal liability; inability to attract and retain talented employees
 - Legal: contract breach; acquisition misstep; compliance issues; loss of authorities
 - Financial: loss of money; personal liability; diminished future appropriations
 - Compliance & Oversight: technology or user inadvertently violates authorities or policies; insufficient documentation of purpose, limitations, accountability
 - Reputational & Trust: damage to standing with customers, overseers, partners, employees, public
 - Transparency: inability to trace AI methodology and algorithms
 - Work product & Explainability: data integrity issues; analysis or other output is biased or incorrect; unanticipated impacts; unable to explain or interpret output
 - Privacy and Civil Liberties: unexpected impact to individuals' privacy or civil liberties

¹⁴⁰ See Appendix C

¹⁴¹ See *Id.*

An initial risk level is determined by combining the likelihood of a failure with the severity of impact across the relevant risk areas.¹⁴² For example, is there is a low, moderate, or high likelihood of supply chain compromise? Would such a compromise affect only one discrete system or would there be system-wide implications? These calculations will result in an initial risk level. Then potential mitigation measures, such as additional policies, training, or security measures, are applied to decrease the initial risk level to an adjusted risk level. For example, physically or logically segmenting an organization's systems so that a compromise only touches one system would significantly reduce the risk level associated with that particular technology. Organizations should seek to apply more than one preventative or mitigative action to further reduce risk and for a more effective layered defense.¹⁴³ The higher the likelihood, the lower the severity must be to offset the risk, and vice versa.

Then organizations must consider that adjusted risk level in relation to their tolerance for risk; how much risk (and potential consequence) is acceptable in pursuit of value? This requires fully understanding the benefits of the proposed activity, opportunity costs of inaction, and the IC's risk tolerance levels. The ODNI should create with IC leadership overarching risk tolerance levels for AI failures in the IC, within which the IC elements may define their own tolerance levels based on their unique missions. Risk tolerance levels will vary for each type of associated activity or technology, but can follow a similar path:

- Mid-level leadership develops common AI and emerging technology business and use cases, including likely risks.
- The RMC develops risk tolerance statements, documenting specific examples of “acceptable” and “unacceptable” risk to help guide the workforce.
 - This is not an exhaustive list but merely examples for the workforce.
 - Organization must update and revalidate examples annually.
- Mid-level leadership communicates risk tolerance internally and externally to stakeholders, and supports information sharing amongst staff to ensure accurate and meaningful risk assessments.

Understanding and considering the risk of action is an important step forward for the IC, but it is not the last step. Sometimes overlooked in risk assessment practices is the consideration of the risk of inaction. To fully evaluate potential options, decision-makers must consider whether the overall risk of doing something is outweighed by the risks of not doing it. If the IC does not pursue particular AI capabilities, what is the opportunity cost of that inaction? Any final determination about whether to take action must consider whether declining to act would cause greater risk of significant harm. While the answer will not always be yes, in the case of AI and emerging technology, it is a very realistic possibility.

Finally, the ultimate decision and associated mitigation steps are only effective if they are clearly communicated up and down the chain and across the multidisciplinary team, ensuring a full understanding by all stakeholders. Broad communication – about the existence of the framework, how to apply it, and expectations for doing so – is vital.

¹⁴² See *Id.*

¹⁴³ The Swiss Cheese Model of Accident Causation is a helpful explanation of the value of layered defenses. See e.g., Jen Stern, *The Swiss cheese model of risk management*, EstateLiving (Mar. 26, 2021), <https://www.estate-living.co.za/news/the-swiss-cheese-model-of-risk-management/>; *Swiss Cheese Model*, U. Mich., http://websites.umich.edu/~safeche/swiss_cheese.html; *The James Reason Swiss Cheese Failure*, What's the Point (May 30, 2018), <https://whatsthepoint.blog/2018/05/30/the-james-reason-swiss-cheese-failure-model-in-300-seconds/>.

This risk framework is designed as an internal tool for use before a problem arises, when IC elements are still considering the development or acquisition of new AI technologies and upgrades, all the way through operation and retirement of the technology. The results of each risk assessment should inform decision-makers' calculation about risks associated with a specific technology, potential mitigation steps, and whether the overall risk is outweighed by the risk of not pursuing a particular technology – the opportunity cost of inaction.¹⁴⁴

The creation of such an IC AI risk framework, with the supporting processes in each IC element, will ultimately create a culture that encourages officers to both proactively consider risk at each stage in the process and solve problems with creativity and innovation because the risk is discussed, understood, and accepted.

EXTERNAL STAKEHOLDERS

An IC-wide AI risk framework will help IC officers determine when and how to take advantage of innovative, but potentially risky, emerging technologies like AI, increasing comfort with uncertainty and risk-taking in the pursuit of new capabilities. Such a risk framework will have even greater impact if it is accepted – explicitly or implicitly – by congressional and executive branch overseers. Although Congress is not formally bound by such a framework, given the significant accountability measures that often flow from these overseers, they should have awareness and an informal ability to provide feedback into the framework as it is being developed.

A meeting of the minds between the IC and its overseers would lead to at least two important benefits: first, increased confidence in the framework across the IC; and second, better insight into IC decision-making for IC overseers. Ultimately, this mutual understanding would encourage exactly what the IC needs to truly take advantage of next-generation technology like AI: a culture of experimentation, innovation, and creativity that sees reasonable risk and failure as necessary steps to game-changing outcomes.

ACTIONS: RISK

1. Create and adopt an IC-wide AI risk framework that is simple, repeatable, and can adapt to changes in the environment. Appendices A - E provide a sample risk framework.
2. Socialize the AI risk framework with congressional and executive branch overseers.
3. Strengthen internal IC processes to manage potential disconnects and create a risk management culture to encourage AI innovation.

¹⁴⁴ Ctr. Strategic & Int'l Stud., *supra* note 38, at 28 ("The PDDNI, working with agency leaders, should develop an innovation risk framework that aids organizations in weighing both the risk of failure/ loss and the opportunity cost of inaction when considering acquisitions and adoptions of innovative new capabilities.").

Meaningful Congressional Oversight – A Critical Component

Congressional oversight of the IC is critical – Congress is the eyes and the ears of the American people and must ensure the IC is wisely spending taxpayer dollars and properly executing national security activities. But intelligence oversight is complicated and has not sufficiently evolved with the times.¹⁴⁵ When it comes to assessing progress of IC programs, standard oversight processes typically track defined, pre-determined requirements, cost, and timelines, which has worked reasonably well for large programs like satellites and fixed facilities like buildings, for which there is a clear beginning, middle, and end, with specified milestones and a definite budget. However such metrics are less effective when it comes to developing AI capabilities, where requirements and timelines are necessarily fluid and the technology is still maturing.¹⁴⁶ Therefore, the priority actions in this report will require corresponding changes in congressional oversight.

The IC's primary congressional oversight committees, the House Permanent Select Committee on Intelligence (HPSCI) and the Senate Select Committee on Intelligence (SSCI),¹⁴⁷ as well as the House Appropriations Committee – Defense (HAC-D) and the Senate Appropriations Committee – Defense (SAC-D), which provide the IC's money, (hereinafter “Committees”), must consider a more adaptive approach to oversight, one that measures progress and failure through metrics that are less rigid and better tailored for AI. In doing so, however, the Committees will lose a measure of certainty that impacts their most powerful lever – fiscal control of the IC. For that reason, the Committees and the IC also must build a greater degree of trust, transparency, and ultimately partnership.

ADAPTIVE OVERSIGHT

Much like AI itself, congressional oversight of AI activities must be more flexible and adaptive than traditional approaches to oversight, to include how Congress measures success. While there are a variety of rules that describe Congress' role in conducting oversight, Congress has considerable latitude and discretion in the execution of that oversight, to include how they measure executive branch progress. Congress and the IC must have a shared strategic vision for what a successful AI project looks like and an approach to oversight that is achievable and tailored for AI.

IC oversight today is most often focused on ensuring projects stay on track in terms of cost and schedule; there are well-defined outputs, such as number of tools built, and static timelines for delivery. Such demonstrable deliverables are objective, consistent, and easy to measure, but incompatible with the nature of AI. AI has made amazing progress in recent years but the underlying technology, the algorithms and their application, are still evolving. AI holds enormous promise to transform a variety of IC missions and tasks, but how and when those changes occur is more difficult to forecast. To take full advantage of AI's emerging possibilities, the IC must have the flexibility to test, evaluate, and adjust as new algorithms and improved capabilities are developed and applied to different problem sets. The IC must experiment and iterate its progress over time. This suggests

145 This is not the first report to call for changes to intelligence oversight; there have been many, including the 9/11 Commission Report. See *CRS Report: Congressional Oversight of Intelligence: Background and Selected Options for Further Reform*, Cong. Rsch. Serv., (Dec. 2018), <https://fas.org/spp/crs/intel/R45421.pdf> (a more detailed discussion of previously proposed significant reforms to intelligence oversight that could greatly benefit the success of the IC).

146 Final Report, *supra* note 7, at 66-68 (The NSCAI final report also identifies the problem with Congress' current approach to oversight and recommends Congress “institute reforms that enable the advancement of software and digital technologies by accounting for speed, uncertainty, experimentation, and continuous upgrades”).

147 The HAC-D and SAC-D appropriate the IC's funding and therefore play a critical role in the IC's ability to execute its mission. To the extent the proposals in this report are relevant to the roles of the HAC-D and SAC-D, the proposals extend to those committees.

there will be unexpected breakthroughs, as well as failures in areas that initially seemed promising; this is a feature, not a bug.

Over the last several years, DOD also has considered changes to oversight, recognizing that detailed requirements, program schedules, milestones, and Gantt charts are not well-suited to measuring factors critical to the success of software.¹⁴⁸ In addition, research by Google and others indicates that certain kinds of metrics, such as those aligned with DevOps¹⁴⁹ best practices, can better predict the performance of software teams.¹⁵⁰ DevOps metrics allow teams to focus on outcomes while adjusting for multi-dimensional, dynamic, and continuous improvement in technology along the way.¹⁵¹ DevOps practices enable software teams to move quickly, respond rapidly to user needs, and produce reliable software – all critical when it comes to adopting AI in the IC.

Therefore, instead of traditional oversight metrics, the IC and the Committees must learn from industry best practices related to DevOps and software, and together develop relevant and adaptive metrics that can be consistently applied but are more aligned with AI's attributes.¹⁵² Performance evaluation and oversight should focus on delivery of capabilities, drilling down on speed and functionality together in phases¹⁵³ and time-boxing segmented activities, from new releases to bug-fixes to staffing. New metrics must focus on key performance indicators that track the progress of how AI tools evolve rather than only the final product, expecting value to the user earlier based on strong communication and feedback loops.

Given that this is a departure from standard metrics and expectations, the IC must explain to the Committees how new metrics will continue to drive accountability and value for these specific projects, and the IC should expect to use new metrics first on low-risk projects. The IC should have the ability to show a project's progress annually – with different types of deliverables – and if timelines slip, the IC must quickly inform the Committees and produce new targets. The Committees must clearly understand the IC's standards and benchmarks so that the Committees can evaluate programs accordingly.

Each oversight conversation should start with a mutual understanding of each others' goals and what success looks like: how does this AI project itself measure value? Increased performance? Reduced cost? Minimized number of humans in the loop? Faster and more comprehensive review of data? Or more likely, some combination of these. New metrics alone will not address whether a programs' goals are worthwhile, so a meeting of the minds between the IC and the Committees on the value proposition of a project is key to measuring progress appropriately; only then should specific metrics be considered.

148 McQuade, *supra* note 45, at S119-23.

149 DevOps and DevSecOps are shorthand for Development and Operations, or Development, Security, and Operations, which encompass various best practices for software development that include rapid updates and tight interdisciplinary collaboration throughout the technology development and deployment process. *DoD Enterprise DevSecOps Reference Design*, Dep't Def. (Sept. 12, 2019), https://dodcio.defense.gov/Portals/0/Documents/DoD%20Enterprise%20DevSecOps%20Reference%20Design%20v1.0_Public%20Release.pdf?ver=2019-09-26-115824-583 ("DevSecOps is an organizational software engineering culture and practice that aims at unifying software development (Dev), security (Sec) and operations (Ops)."); *See generally* *DevSecOps*, Comput. Sec. Res. Ctr. (Oct. 21, 2020), <https://csrc.nist.gov/Projects/devsecops>.

150 Forsgren, N., Humble, J., & Kim, G., *Accelerate: The science of lean software and devops: Building and scaling high performing technology organizations*, It Revolution Press (2018).

151 This is more aligned to a "capabilities model" of measurement rather than a "maturity model."

152 The HPSCI has already recognized this issue, as is evident in its Report that accompanies the FY21 Intelligence Authorization Act. *See generally* H.Rept. No. 116-565 (2020) <https://www.congress.gov/congressional-report/116th-congress/house-report/565> (Metrics for Modern Best Practices in NGA Software Programs).

153 McQuade, *supra* note 45, at S119-23.

Based on DevOps and other agile environments, there are a variety of metrics that could be used for AI projects. Aligning them with the “cost, schedule and performance” trio that is standard for program management,¹⁵⁴ as done below, provides a known and understandable framework for the Committees, as well as IC officers. The following is a non-exhaustive list of metrics that would better indicate the progress and success of AI projects.¹⁵⁵

COST-RELATED METRICS:

- Personnel hours saved
 - Minimized number of humans in loop
- Are program managers aggressively pivoting away from efforts that are not working to new initiatives? Start small, iterate fast, terminate quickly.

SCHEDULE-RELATED METRICS:

- Time from identification of need to contract completion
- Time from program launch to deployment of simplest useful functionality
- Speed of data preparation and conditioning
- Speed of model design, training, testing, deployment, and retraining, if needed
- Deployment frequency for major functions and updates
 - Change fail percentage
- Time from new user-requested feature to deployment
 - Speed of iteration and feedback loops between technologists, operators, analysts, and prototype demos in field

PERFORMANCE:

- Model robustness against adversarial attacks
- Compute costs to operate in the field
- Integration of technology providers and acquisition officers into user experimentation, modeling, simulations to help develop operational concepts that incorporate AI
- User satisfaction
- Faster, more comprehensive results
- Accuracy, repeatability, and consistency of solution

¹⁵⁴ *Program Management*, AcqNotes (last updated May 3, 2021), <https://acqnotes.com/acqnote/careerfields/program-management-overview>.

¹⁵⁵ McQuade, *supra* note 45, at S82

There also must be ethics related metrics, for example:

- Diversity, density, and homogeneity of data and data sources when training and validating the technology
- Level and type of human interaction with the technology
- Explainability and transparency of the technology through:
 - Simplicity of explanations regarding AI decisions
 - Clarity to user that AI is driving the output
 - Predictability of the probable action of the technology
- Robustness of AI governance and security processes

While pivoting to new metrics is a good start, the IC and the Committees both must be open to iterative and adaptive processes. Oversight must have the ability to flex and evolve if the initial approach is less than optimal. A true partnership between the IC and the Committees should allow for healthy conversations about what is and is not working well, so that congressional oversight becomes best tuned for AI.

TRUST, TRANSPARENCY, AND PARTNERSHIP

There is no dearth of oversight today – current processes consist of daily, sometimes hourly, engagement between the Committees and the IC. Each year, there are hundreds of written reports, in-person briefings, phone calls, hearings, and other engagements between congressional overseers and the IC.¹⁵⁶ However, current Committee engagements with the IC suggest a lack of confidence and trust in the IC's activities; these engagements are often excessively tactical and focused on execution details that provide neither a strategic perspective on the health of the program nor an understanding of potential long-term opportunities and risks.

Such a focus is frequently ineffective in that it neither provides the Committees with the insight they need nor leads to the congressional support the IC expects,¹⁵⁷ steering the relationship toward one that is more adversarial than collaborative. Moreover, these engagements drive a continuing cycle of meetings and briefings, requesting deeper levels of detail, in an effort to achieve the desired understanding. Unfortunately, layering detail on top of detail does not produce strategic insight.

To be clear, the current oversight processes were not designed to be overly burdensome or roadblocks to progress – they were designed to give Congress appropriate insight and confidence that executive branch activities and spending are being carried out efficiently, effectively, and consistently with the law. Unfortunately, the processes have become onerous due to a history of issues that have understandably undermined Congress' trust and confidence in the IC.¹⁵⁸ The IC must rebuild trust with Congress so that Congress can step back from day-to-day operational details and oversee the Community at a more appropriate strategic level.

¹⁵⁶ Rightly Scaled, *supra* note 35, at 33.

¹⁵⁷ *Id.*

¹⁵⁸ CRS Report, *supra* note 145, at 3-4 (describing history of intelligence oversight committees); See also Loch K. Johnson, *The Contemporary Presidency: Presidents, Lawmakers, and Spies: Intelligence Accountability in the United States*, 34 Presidential Stud. Q. 828, 829-833 (2004) (describing the history and evolution of the relationship between Congress and the IC).

Consider the relationship between a Board of Directors (Board) and a Chief Executive Officer (CEO) in the private sector.¹⁵⁹ The Board has ultimate responsibility for overseeing the affairs of the organization and ensuring the organization is appropriately stewarding the resources entrusted to it. Managing the execution of a company's day-to-day activities is the job of the CEO.

According to the KPMG Board Leadership Center, the key to a healthy relationship between a Board and the organization it oversees is trust and transparency, where the Board has constructive conversations with the leadership team about significant decisions and issues, the Board has the opportunity to provide meaningful input before decisions are made, and the leadership team receives feedback it considers to be valuable.¹⁶⁰ It is not the Board's role to "see every scrap of paper that the management team sees."¹⁶¹ The Board should not wade into the tactical details of an issue unless it is strategy or risk related. Instead, Board members must make the effort to know about the company and apply their wisdom from analogous situations and decades of experience to execute rigorous oversight, without seeing every last spreadsheet. This requires a mutual understanding of each others' goals and an acceptance of each others' roles.

This analogy is useful in thinking about the oversight relationship between Congress and the IC, where the DNI is the CEO, and the Committees are the Board. The Committees have the responsibility to ensure the IC is appropriately stewarding its resources, and in doing so they also have the responsibility to leverage their knowledge and expertise to provide strategic advice and direction rather than diving into the many levels of detail below. But, as with the private sector, without the requisite trust and transparency it is difficult for the Committees to operate at the strategic level.

To rebuild the needed trust, the IC and the Committees must fundamentally change the nature of their interaction and significantly rethink the communication and transparency between them.¹⁶² They must be more open with each other and regularly engage as partners in formal and informal conversations about AI activities to better manage expectations and reduce the element of surprise. The IC should seek Committee views on significant activities before final decisions are made and work to incorporate their feedback when possible to build stronger support and buy-in from the Committees.¹⁶³ The Committees must have confidence that the IC will engage early and with reasonable solutions when something goes wrong. And in return the IC must know that the Committees will leverage their expertise to provide strategic insights and drive supportive

159 See also Bilal Zuberi and Anthony Thomas, *It's Time For the DOD and the Emerging Tech Sector to Make 'Big Bets' on Each Other: Four Ways to Sharpen Our Technological Edge in Defense* (Sept. 2020), <https://www.linkedin.com/pulse/its-time-dod-emerging-tech-sector-make-big-bets-each-other-zuberi/> (noting that Congress should be treated like DOD's Board of Directors).

160 *Insight: Is The Board Helping The Business Keep Pace?*, KPMG Bd. Leadership Ctr. 1, 1-2, <https://boardleadership.kpmg.us/relevant-topics/articles/2021/is-the-board-helping-the-business-keep-pace.html> ("Rather than just management reporting up to the board or the board making pronouncements on high, it has to be a conversation... informed by trust, where management feels like they're going to get useful feedback from the directors and the directors feel like they're not just being presented to on things that are fait accompli, but they're being asked to provide useful input.").

161 *Id.*

162 Johnson, *supra* note 159, at 834 (explaining that oversight works only if there is honesty and completeness in the relationship between Congress and the IC).

163 Presidents have asserted constitutional authority to withhold certain information from Congress for decades, to include executive branch deliberations – that is, information that is part of the internal executive branch deliberative process that leads to a final decision. See Congressional Requests for Confidential Executive Branch Information, Op. O.L.C. 153, 155 (1989). This causes tension with the Committees as they must oversee programs they had minimal insight or input into before they were created, and may fundamentally disagree with or have extensive experience that might have benefitted the program in advance, leading to a more adversarial relationship than one of partnership. If the deliberative process privilege becomes a problem, the IC could use the well-established, if formal, "accommodation process," which allows the executive branch to provide such information as is necessary to accommodate the legitimate needs of the legislative branch, to provide earlier insights to the Committees on emerging AI programs and issues. See *Id.* at 157-161. This kind of accommodation would not only bolster trust amongst the IC and the Committees, it also would provide an opportunity for the Committees to give helpful feedback and input to the IC, and it would inform future AI legislation.

legislation and funding – even in times of loss or failure – and that the Committees will publicly back the strategic direction of the IC, even in the face of criticism.

A major mindset shift is required to make this happen, and that mindset shift starts at the top. The DNI and IC leadership must model forward-leaning behavior that treats the Committees as the critical partners they are. The DNI must clearly set expectations with Committee leadership, and then promote a new approach throughout the IC. The approach should include proactive, standing formal and informal IC leadership engagements with the Committees, as well as IC legislative engagement principles, promulgated by the DNI, that set forth the DNI's expectations for IC formal and informal discussions with the Committees.

Leaning forward in this way does come with some risk to the IC: with more information, the Committees can do more to stop projects they may not like or to steer projects in certain directions. To truly build trust, the Committees must be judicious in these engagements, focusing on insightful strategic and risk-based questions reflective of their extensive experience and expertise, and they must not misuse the information to inappropriately interfere with the executive branch's authority to execute the law. The following steps would encourage a positive shift.

First, to increase their capacity and institutional expertise, the Committees should consider reorganizing staff along functional lines, rather than by IC element. By doing so, Committee staff would develop a greater understanding of various AI tools and technologies that could be applied strategically across IC elements and activities, and they would have a more holistic cross-IC understanding of AI coordination and activities. While the more common model of organizing staff by IC element makes logical sense, expecting staff to understand the entirety of what an IC element does – from operations, to capabilities, to recruiting and retention, to public private partnership, and everything in between – is unrealistic. This is exceptionally difficult even for IC leadership to do and is unreasonable to expect of Committee staff who are often single-threaded in their roles. Refocusing staff on specific functional areas and allowing them to become experts would greatly benefit not only the Committees' oversight of those activities, but the IC elements they oversee. The HPSCI has recently done just this;¹⁶⁴ the IC's other oversight committees should consider a similar change. In addition, Congress more broadly should revive or recreate the Office of Technology Assessment to allow the Committees the ability to consult technical experts with deep background in AI and related technologies, when needed.¹⁶⁵¹⁶⁶

Second, IC elements should propose two informal engagements¹⁶⁷ with the Committees: semi-annual conversations between Committee staff and high-priority AI project leads for conversation and feedback on progress, issues, concerns, and requirements; and periodic DNI and IC leadership "coffee catch-ups" with Committee members to better drive the strategic relationship, ensure each has the benefit of each other's thinking at that moment, and develop a sense of partnership. These informal engagements should not follow scripts or seek to accomplish specific tasks, but rather should be focused on creating shared understandings, open dialogue, and trust-building around AI

164 Rightly Scaled, *supra* note 35, at 22.

165 *Id.* at 35; See also *Report of the Task Force Project on Congressional Reform*, Am. Political Sci. Ass'n 1, 8 (Oct. 29, 2019), <https://www.legbranch.org/app/uploads/2019/11/APSA-Congressional-Reform-Task-Force-Report-11-2019-1.pdf> (discussing the value of reviving or recreating the Office of Technology Assessment to add general-use, neutral, fact-based expertise to Congress).

166 *Final Report*, Cyber Solarium Comm'n 1, 36 (Mar. 2020), <https://www.solarium.gov/>.

167 In order to promote trust and candor, both the informal and formal engagements should be closed to the public and press, regardless of classification level.

activities. The AI project leads should consistently reiterate what is known and unknown about the project, all of the potential outcomes, and any changes in spending the Committee may see in the coming months. The senior leadership coffee catch-ups would, of course, produce benefits well beyond the IC's AI activities.

Third, if AI receives no-year or multi-year funding, the Committees should hold a focused annual review of AI spending during the previous year. This review should include sufficient detail to provide the Committees an understanding of what is going well and what did not go as expected, so the Committees can provide a timely and critical check on the use of that money. If the funding has been executed in accordance with Congressional direction – even if some of the activities have failed – the money should continue to flow. If the funding has not been executed properly or consistently with Congressional direction, Congress should have the ability to stop the funding immediately.

Fourth, the Committees should request formal semi-annual substantive briefings from the IC on high-priority AI projects, as defined jointly each year. These briefings should include a level of detail that allows the Committees and staff to understand progress against the new metrics, and ask questions about the strategic direction of the programs, areas of risk, concerns, unexpected issues, and future needs in terms of legislation or funding. Ultimately, these substantive briefings should provide both the IC a mechanism to show forward movement and elevate significant issues, and the Committees a mechanism to more formally track high-priority AI activities across the IC.

The Committees, just like a corporate board, provide an important check and balance on the IC's AI activities. But to be successful, the Committees' involvement must be strategic and focused on the health of the overall project, leaving the tactical execution of the activities to the IC. This requires trust, transparency and partnership, with both the IC and the Committees coming to terms with their respective roles and tools to influence outcomes. This kind of approach would very likely take some practice to get right but, if successful, could dramatically change the IC's partnership with the Committees for the better, providing earlier and increased insights to the Committees and leading to greater support and backing for the IC.

ACTIONS: OVERSIGHT

1. The IC should create new, AI-specific metrics, in consultation with the Committees.
2. The Committees should reorganize staff portfolios into functional areas and Congress should revive the Office of Technology Assessment to improve Congress' technical literacy.¹⁶⁸¹⁶⁹
3. The IC should adopt a more informal and transparent posture with Congress.
 - a. The DNI should define IC legislative engagement principles to clearly set the tone and guidelines for formal and informal engagements with the Committees.
 - b. AI project leads should hold semi-annual informal conversations Committee staff.
 - c. IC leadership should hold periodic informal Coffee Catch-ups with Committee members.
4. If the IC receives no-year or multi-year funding, the Committees should create a formal annual AI budget assessment activity.
5. The Committees should hold formal semi-annual strategic AI briefings to review agreed upon AI metrics and engage on issues, concerns, and needed funding or legislation.

¹⁶⁸ Rightly Scaled, *supra* note 35, at 35.

¹⁶⁹ Cyber Solarium Comm'n., *supra* note 167, at 36.

PART III: SUMMARY OF ACTIONS

CULTURE OF INNOVATION:

1. The DNI must clearly and publicly prioritize innovation. The DNI and IC leadership should:
 - a. Designate a senior leader to drive innovation across the IC as a top priority.
 - b. Set clear goals and expectations for IC innovation activities.
 - c. Articulate acceptable risk and expected failure rates.
 - d. Change incentive structures to reward creativity and innovation.
2. The DNI should require a 90-day action plan that takes advantage of the AI burning platform to drive a series of proposals and activities that will begin to turn the IC culture toward one that embraces and supports innovation broadly.
 - a. These proposals should articulate expectations for how innovation will be integrated across IC activities, rather than done in independent stovepipes.

NATIONAL SECURITY ECOSYSTEM:

1. ODNI and DOD General Counsels and Legislative Affairs Chiefs should meet semi-annually to discuss proposed legislation that may be applicable to the broader IC and DOD.
 - a. They also should create a specific, formal step in the legislative process to ask whether the other might need similar legislation.
2. The ODNI and DOD should instruct all IC and DOD officers, with specific attention to legal and legislative officers, to proactively consider who else may need relevant authorities and propose adjustments, if appropriate. This includes ensuring subject matter experts are consulted to understand the objectives of the authorities and programs.
 - a. This is doubly important for the ODNI team, as they have a unique cross-IC optic and awareness and it is specifically their mission to look out for the entire community.
3. The NSCAI's Steering Committee on Emerging Technology should require annual briefings from the ODNI and DOD General Counsels and Legislative Affairs Directors on their efforts to better coordinate AI-related pilots, authorities, and proposals.

BUDGET & ACQUISITION (FLEXIBILITY & SPEED):

1. The ODNI should require certain types of programs to use the SOO/PWS model by default.
 - a. Identify criteria for programs that should use an SOO/PWS rather than an SOW.
 - b. Identify SOO/PWS experts in each organization who can mentor less experienced officers and identify appropriate additional training.
 - c. Accelerate creation of an AAF that includes a software acquisition path, learning from DOD's experiences.
2. Congress should immediately provide OT and CSO authority to the IC.
3. The ODNI should create an unclassified, stand-alone IC sandbox housed in ICWERX where IC customers can test and evaluate AI capabilities quickly to buy down risk and speed up acquisition.
 - a. Congress should authorize the IC to quickly and easily spend for T&E purposes up to \$100,000 per capability to facilitate the sandbox evaluation process.
4. Congress should extend DOD's single appropriation for software pilot to the IC.
5. The ODNI should work with Congress to create no-year funds for AI.
 - a. At a minimum, all AI must be designated as part of RDT&E 2-year funds.
6. The ODNI should work with Congress to establish an IC AI Technology Fund, administered by the DNI's Director of Science & Technology.
 - a. Must have appropriated funds that are no-year, no-color, and non-taxable.
 - b. Define what qualifies as "AI" for purposes of this fund.¹⁷⁰
7. The ODNI should use the AIM iHub as the foundation for an ICWERX, modeled after AFWERX, to increase speed, agility, and ease of engagement with private sector innovators and small businesses developing AI tools and technologies.
 - a. Create a loose partnership with AFWERX, to learn from their experts and leverage their strengths while building the expertise of the IC's acquisition professionals regarding the use of authorities and market research.
 - i. Consider co-location of small ICWERX within AFWERX to start.
 - ii. Prioritize and hire additional acquisition and market research experts (see below).
 - b. Leverage and increase iHub/ICWERX connections to other IC innovation labs and centers to share expertise and consult. This will enable a better understanding of the IC's capabilities and needs, better training of the IC's workforce, better use of IC authorities, better partnering with the private sector, and ultimately faster time to mission for emerging technologies.
 - c. Use AI Technology Fund to fund specific activities within ICWERX.
 - d. House unclassified IC sandbox in ICWERX for quick and easy T&E of new AI capabilities.

¹⁷⁰ The DNI's AIM strategy, which defines AI, should be the starting point. See generally *The Aim Initiative*, Director of National Intelligence, <https://www.dni.gov/files/ODNI/documents/AIM-Strategy.pdf>.

BUDGET & ACQUISITION (COMPLEXITY):

1. The IC must fully understand and use flexible authorities it already has for both budget and acquisition.
 - a. Partner with and leverage USG-wide expertise (like TTS and FEDSIM) to increase IC's knowledge base.
 - i. Strengthen partnership with the DOD JAIC acquisition team, including engagement on its Tradewind initiative.
 - b. Develop comprehensive IC training and development program for everyone from acquisition, contract and procurement officers to lawyers to operators and others associated with the acquisition. The ODNI should:
 - i. Partner with existing programs at DIA, DHS, DOD and elsewhere to enable IC officers to participate in existing courses.
 - ii. Charge NIU with helping to address the IC education gap in acquisition and emerging technology through workshops, a certificate program, or by linking courses at DIA, DHS, DOD and elsewhere.
 - iii. Create courses for various levels of officers, including managers and supervisors, who should understand acquisition requirements and processes at a high level and the impact that work has on the mission.
 - iv. Ensure training explores existing FAR flexibilities, OTA, SBIR/STTR, Statements of Objectives, market research, and data rights.
 - v. Ensure contractors working on program management or acquisition activities have sufficient expertise and meet training standards.
 - vi. Strengthen continuing education requirements for all IC acquisition professionals to keep skills fresh and reward those who stay current on critical skills with incentive pay.
 - c. Prioritize and speed the recruitment and hiring of highly qualified acquisition professionals using existing direct hire authorities.
 - d. Further develop and accelerate promulgation of the IC Acquisition Playbook and lessons learned to every IC acquisition professional to both ensure the ability to use authorities and to encourage a community approach to usage.
 - i. Start by posting the current Playbook on the ODNI website so it is easily accessible to IC officers.
 - ii. Link IC Playbook to FAI Periodic Table for additional authorities and examples.
2. The ODNI should align incentives with desired behavior.
 - a. Create a collaborative team approach that brings acquisition professionals into projects earlier in the process.

- b. Evaluate performance of IC acquisition professionals, their supervisors and managers against processes and outcomes, in addition to funding obligation rates.
 - i. Include focus on engagement and collaboration with customers and partners, meaningful decision processes, delivering mission outcomes, and learning from failures.
 - ii. Require project teams to engage with acquisition professionals early.
- c. Reinvigorate IC-wide acquisition and mission awards that recognize acquisition professionals' mission impact, innovation, and cross-community collaboration.

RISK:

- 1. Create and adopt an IC-wide AI risk framework that is simple, repeatable, and can adapt to changes in the environment. Appendices A - E provide a sample risk framework.
- 2. Socialize the AI risk framework with congressional and executive branch overseers.
- 3. Strengthen internal IC processes to manage potential disconnects and create a risk management culture to encourage AI innovation.

OVERSIGHT:

- 1. The IC should create new, AI-specific metrics, in consultation with the Committees.
- 2. The Committees should reorganize staff portfolios into functional areas and Congress should revive the Office of Technology Assessment to improve Congress' technical literacy.
- 3. The IC should adopt a more informal and transparent posture with Congress.
 - a. The DNI should define IC legislative engagement principles to clearly set the tone and guidelines for formal and informal engagements with the Committees.
 - b. AI project leads should hold semi-annual informal conversations Committee staff.
 - c. IC leadership should hold periodic informal Coffee Catch-ups with Committee members.
- 4. If the IC receives no-year or multi-year funding, the Committees should create a formal annual AI budget assessment activity.
- 5. The Committees should hold formal semi-annual strategic AI briefings to review agreed upon AI metrics and engage on issues, concerns, and needed funding or legislation.

APPENDIX A: RISK FRAMEWORK INTRODUCTION¹⁷¹

The IC should create common, overarching AI risk management principles (“IC Principles”), agreed to by the directors of each IC element. These IC Principles should articulate IC definitions, thresholds, tolerance levels, and standards, which could be drawn from Appendices B-E below. The IC Principles must drive a repeatable risk assessment framework (below), to help officers assess potential AI failures and associated risks, and ensure thorough internal consideration of issues in advance of final decisions. This IC AI Risk assessment framework has five major steps:

- Failure Identification (identify potential failures and range of effects)
- Risk Analysis (analyze failures and risks)
- Risk Assessment (prioritize and reconcile risks, mitigations)
- Decision-making (develop action plan, communicate)
- Monitor and update (evaluate results for further action)

This framework should be used throughout the lifecycle of the activity being assessed. The AI and emerging technology lifecycle of Plan, Design, Build/Buy, Test, Deploy, Operate, Maintain, Decommission includes the following activities during which risk should be considered/assessed:

- Identify technology/private sector entity
- Engage private sector entity
- Build/acquire technology, test/modify/adapt as necessary
- Secure technology
- Use technology
- Assess technology performance
- Update technology
- Disposition technology

Risk assessments should be completed by officers seeking to acquire or use new AI tools and technology that meet the IC element’s threshold to complete a risk assessment. Each IC element’s Risk Management Council should help set its thresholds, which must be consistent with the IC Principles and socialized across the community to level-set expectations and support interoperability.

¹⁷¹ Appendices A-E were drawn from the following works: *Guide for Conducting Risk Assessments*, Nat’l Inst. Standards & Tech. 1, D1-H1 (Sept. 2012), <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf> (framework was used as an example or model for Tables 1, 2, 3, 5, and 6); *AI Using Standards to Mitigate Risks*, Dept. Homeland Sec. (2018), https://www.dhs.gov/sites/default/files/publications/2018/AEP_Artificial_Intelligence.pdf; *Report on Algorithmic Risk Assessment Tools in the U.S. Criminal Justice System*, Partnership on AI, <https://www.partnershiponai.org/wp-content/uploads/2019/04/Report-on-Algorithmic-Risk-Assessment-Tools.pdf>; Mark Latonero, *Governing Artificial Intelligence: Upholding Human Rights & Dignity*, Data & Society, https://datasociety.net/wp-content/uploads/2018/10/DataSociety_Governing_Artificial_Intelligence_Upholding_Human_Rights.pdf; Raji et al., *Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing*, Fairness, Accountability, and Transparency Conference 33, 36 (Jan. 27-30, 2020), <https://dl.acm.org/doi/epdf/10.1145/3351095.3372873>.

Thresholds may be based on dollar value, scope of access to systems, privacy implications, or other issues that impact the range of effects and whether they would be sweeping, wide-ranging, extensive, limited, or minimal.

Risk assessments require a multidisciplinary approach, with significant input from factual and subject matter experts across the spectrum of relevant functions (collectors, analysts, attorneys, acquisition specialists, security officers, privacy officers, etc.). In addition, officers completing risk assessments must regularly consult with their supervisors and managers to ensure proper levels of awareness and elevation to senior levels, as necessary.

APPENDIX B: PREPARING THE FRAMEWORK

Process and business drivers

Before the risk framework is used within an organization, that organization must ensure proper governance and **internal processes** are in place, to include an multidisciplinary **Risk Management Council**, and articulate the relevant **business drivers**. The organization's business drivers should nest within the IC's business drivers and be tailored to the organization's specific mission.

Examples of IC business drivers related to AI tools include:

- The ability to better analyze and understand voluminous data in a timely fashion, which supports:
 - Policymaker decision cycle
 - Intelligence targeting
 - Intelligence collection
 - Intelligence analysis
 - Detection of adversary tactics, techniques, and procedures targeting the US
 - Detection of insider threats
 - IC business analytics
 - IC enabling activities
- The ability to remain competitive with adversaries using similar technologies
- The ability to attract and retain a highly technical and talented workforce
- The ability to help shape and lead global AI norms, ethics and standards
- The need to carefully use taxpayer dollars
- The need to ensure the security of IC systems, information, and people
- The need to retain authorities

Risk Tolerance and Decision-makers

Every organization must take risk to achieve its objectives and sometimes the greatest risk is inaction. Therefore, each IC element must determine the level of risk it is willing to accept and where decisions will lie. IC element risk tolerance statements must be consistent with the IC Principles. Clarity on an organization's **risk tolerance** will help limit excessive risk-taking as well as unnecessary risk aversion, both of which can impede progress and success. It will encourage IC officers and stakeholders to take agreed and calculated risks, to recognize uncertainty, to plan for the possibility of failure, and to learn from both positive and negative results. As a result, an organization might mitigate, transfer, avoid, or accept risk depending on the potential impact to the delivery of critical services. Examples of risk tolerance statements include:

- The IC will pursue AI technologies with a commercial partner only when those technologies are pursuant to an agreement or strategy and approved by appropriate decision levels.
 - Acceptable risk (where mission value exceeds risk): the IC pursues technologies designed for foreign intelligence purposes and consistent with documented IC element priorities with center or directorate level approval.
 - Unacceptable risk (where risk exceeds mission value): the IC pursues technologies designed for undefined purposes without a clear understanding of how it will further IC element priorities without center or directorate level approval.
- The IC will retain and use large scale data with sufficient civil liberties and privacy safeguards to reduce impact on individuals and entities.
 - Acceptable risk: The IC uses data that has been reviewed and approved by the element's Civil Liberties Protection Officer with center or directorate level approval.
 - Unacceptable risk: The IC uses data without approval by the element's CLPO or IC element head.
- The IC will endeavor to use AI tools that include the ability to exit from an AI vendor or solution to retain flexibility to respond to changes in the environment. Use of tools that are tied into a vendor or solution will be used only where there is the potential for high value outcome that justifies the risk.
- The IC will not transfer risk to another organization unless the associated impacts have been considered and communicated to customers, when feasible, and a mitigation plan has been put in place, or it has been approved by the IC element head or deputy.
- The IC will not procure new AI tools and technologies unless it has sufficient talent to use the tools or develops a talent development plan in parallel.
- The IC will endeavor to use AI tools that have functionality to explain how they came to certain outcomes. Use of tools with less explainability will be used only when there is the potential for high value outcome that justifies the risk and such use is approved at the center or directorate level.
- The IC will only use AI tools if adequate security, monitoring and oversight processes exist to ensure appropriate implementation and solutions.

Another approach to risk tolerance statements is:

- The IC in general has a high tolerance for risks required to introduce and scale technology that will help solve priority 1 intelligence problems (and a low tolerance if solving only priority 4 intelligence problems).
- The IC in general has a low to tolerance for risks associated with US persons' privacy and civil liberties.
- The IC seeks to reduce the risks associated with the privacy and civil liberties of non-US persons.
- The IC in general has a low tolerance for risks that would result in the degradation of US public confidence in the IC.

IC elements also must clearly articulate where **decisions and accountability** will lie. For example, for “very low” risk activities, decisions might be made at the project level. For “low” risk activities, decisions might be made at the office level. For “moderate” risk activities, decisions might be made at the group level (which is higher than office). For “high” risk activities, decisions might be made at the directorate level. For “very high” risk activities, decisions might be made at the IC element director level. Articulating this clearly in advance will help everyone understand expectations as their activity progresses.

APPENDIX C: USING THE FRAMEWORK

Table 1 provides the general categories and activities that will be required as part of each risk assessment. Before starting the risk assessment, ensure a clear understanding of overall purpose, scope, assumptions, constraints, and information sources for risk assessment. Then follow these steps:

1. Identify potential failures and range of effects for specific AI tool or technology (Tables 2 & 3)
2. Analyze failures and risks by identifying organization's vulnerabilities
3. Assess, prioritize, and reconcile risks by evaluating likelihood and severity of potential impacts
 - a. Calculate initial risk level by combining the likelihood (Table 5) of impact with the severity (Table 6) of impact
 - b. Identify and assess potential mitigations (Table 4)
 - c. Re-calculate the level of adjusted risk factoring in mitigations (Table 7)
4. Decide whether to mitigate, transfer, avoid, or accept the risk
 - a. Balance adjusted risk level against risk of inaction and opportunity costs
 - b. Create an action plan
 - c. Communicate and share risk assessment results
5. Maintain, monitor, and periodically update the risk assessment; Evaluate results on ongoing basis

TABLE 1: GENERAL RISK ASSESSMENT ACTIVITIES

<i>Categories</i>	<i>Actions</i>
Potential Failures	Identify potential failures and sources through the first order and second order risks.
Range of Effects	Identify the range of effects from the failure (sweeping, extensive, wide-ranging, limited, minimal), considering vulnerabilities and mitigations.
Vulnerabilities	Identify existing or potential weaknesses that could enable or exacerbate the failure and any predisposing conditions, such as classified information, that could increase the likelihood of adverse impacts.
Likelihood	Determine the likelihood that failure will be initiated and have adverse impact over a certain time period, considering sources, vulnerabilities, and mitigations.
Severity	Determine the severity of the adverse impact (i.e., potential harm to organizational operations, organizational assets, individuals, other organizations, or the Nation) from the failure. Consider also immediate and potential future impact.
Mitigations	Identify any technical, policy, security or other mitigations, safeguards or counter-measures that may help blunt likelihood of occurrence or severity of adverse impact.
Risk Levels	Determine the ultimate level of risk as a combination of likelihood of occurrence and severity of impact. Then determine adjusted overall risk by factoring in mitigations.

TABLE 2: EXAMPLES OF FIRST ORDER RISKS (RISKS THAT WILL START A CHAIN OF EVENTS THAT LEAD TO IMPACT)

<i>AI Initial Risks</i>	<i>Potential Sources of Failure</i>
Technology	<p>Development, testing, or operational performance cannot be completed as expected or does not work in practice.</p> <p>Infrastructure failures, including failures of equipment, environmental controls, or software due to aging, resource depletion, or other circumstances which exceed expected operating parameters.</p> <p>Implementation errors, incorrect usage of technology</p> <p>Requirements shift or change significantly</p>
Data & Algorithms	<p>Insufficient, incorrect, biased, unsecured, or other faulty data inputs.</p> <p>Incorrect storage, access, retention, dissemination, disposition of data.</p> <p>Algorithms that include insufficient, incorrect, biased or other faulty instructions.</p>
Intentional adversarial attack	<p>Adversarial actor has capability and intentionally targets AI capabilities to deny, disrupt, degrade or destroy capabilities or to infiltrate organization.</p>
Security	<p>System is not sufficiently secure, resilient or redundant, allowing for internal individuals to misuse their access to deny, disrupt, degrade, or destroy capability or external individuals to infiltrate for same purposes.</p> <p>New technology could be leaked, stolen, or weaponized.</p>
Supply chain security or integrity breach	<p>AI component parts are intentionally or unintentionally damaged or compromised such that technology does not work or cannot be trusted.</p> <p>Technology, or components, comes from untrusted country/origin</p>
Operational	<p>Human error in operating technology</p> <p>Used for inappropriate or overbroad purpose</p>
Partner	<p>Company or supplier has failure that impacts ability to provide quality, reliable inputs.</p> <p>Foreign ownership or stake</p> <p>Supply chain threat</p>
Funding	<p>Insufficient funds or cost overruns.</p>
Schedule	<p>Delays render technology less useful or obsolete.</p>
Policy	<p>Unclear or insufficient standards, policies, or processes to properly manage and evaluate the data, technology, usage, or other related activities.</p>

TABLE 3: EXAMPLES OF SECOND ORDER RISKS (RISKS THAT COULD RESULT FROM THE FIRST ORDER RISKS AND LEAD TO COMPOUNDED IMPACT)

<i>AI Derived Risks</i>	<i>Potential Sources of Failure</i>
Workforce	Workforce skills do not match technology needs, insufficient training Personal liability for contracting officers and others who commit funds Inability to attract and retain talented employees
Legal	Breach of contract; acquisition misstep; compliance issues; loss of authorities, ethics issues
Financial	Loss of money; personal liability for contractual obligations; diminished future appropriations
Compliance & Oversight	Technology or user inadvertently violates authorities or policies Accountability of users not clear or enforceable Insufficient documentation of purpose, limitations, accountability; approaches to user-based collaboration and information sharing
Reputation & Trust	Damage to agency reputation, standing with customers, overseers, partners, employees, public Damage to partner reputation, survival of company
Transparency	Inability to understand or trace AI methodology and algorithms
Work Product & Explainability	Data integrity issues; analysis or other output is biased or incorrect; unanticipated impacts; unable to explain or interpret output
Privacy & Civil Liberties	Adverse impact to individuals' privacy and/or civil liberties, including implications for individual personally identifiable information (PII), equal protection, freedom of expression or due process; human rights or international humanitarian rights implications

TABLE 4: POTENTIAL MITIGATIONS

<i>Categories</i>	<i>Potential Mitigations</i>
Technical & Industry	<ul style="list-style-type: none"> • Design Checklists • Failure Modes and Effects Analysis; failover processes; protective technology • Alternative or approved suppliers • Vendor component inventory • Network segmentation • Off-ramps and down-select options • Technical Security Criteria, Standards and Assessments • Documentation of connections between and dependencies within systems • Robust test and evaluation processes • Third party testing centers (National Labs, FFRDCs, UARCs) for high-risk systems • Documentation of required maintenance, re-testing, and technical refresh • Measures of effectiveness • Technical Audit and Monitoring
Data & Information Security	<ul style="list-style-type: none"> • Documented standards for collection, protection, sharing or use of relevant data • Documented origin and provenance of data, outputs, and test results. • Identity and Access Management Controls • Periodic Security Assessments • Active Red Team for both intentional and unintentional failures • Proactive vulnerability assessment, documentation, and remediation
Legal & Regulatory	<ul style="list-style-type: none"> • Clear and Documented Usage Rules and Governance, Roles and Responsibilities • Documented Risks, Appropriate Use Cases, Mitigations • More Targeted Usage or Narrowed Purpose • Marking Outputs as Derived from AI • Training • Oversight Team and Processes • Policies and Technology to Strengthen Accountability • Processes and Infrastructure for Traceability, Audit and Monitoring • Crisis Management Plan • Lessons Learned and Feedback Processes
Moral & Ethical	<ul style="list-style-type: none"> • Transparency • Standards, Governance, and Processes to Discover and Address Bias Throughout Lifecycle • Metrics to Assess AI Output Accuracy • User and Peer Engagement in Model Development and Review Process • Independent Review of AI Tool, Purpose, Intended Use • Privacy Impact Assessment during pre-deployment phase

TABLE 5: DEFINING RISK LEVELS - LIKELIHOOD

<i>Qualitative Values</i>	<i>Qualitative Values</i>		<i>Description</i>
Very High	96-100	10	Failure is almost certain to occur.
High	80-95	8	Failure is highly likely to occur.
Moderate	21-79	5	Failure is somewhat likely to occur.
Low	5-20	2	Failure is unlikely to occur.
Very Low	0-4	0	Failure is highly unlikely to occur.

TABLE 6: DEFINING RISK LEVELS - SEVERITY

<i>Qualitative Values</i>	<i>Qualitative Values</i>		<i>Description</i>
Very High	96-100	10	The failure could be expected to have multiple severe or catastrophic adverse effects on organizational operations, organizational assets, individuals, other organizations, or the Nation.
High	76-95	8	The failure could be expected to have a severe or catastrophic adverse effect on organizational operations, organizational assets, individuals, other organizations, or the Nation. A severe or catastrophic adverse effect means that, for example, the failure might: (1) cause a severe degradation in or loss of mission capability to an extent and duration that the organization is not able to perform one or more of its primary functions; (2) result in major damage to organizational assets; (3) result in major financial loss; or (4) result in severe or catastrophic harm to individuals involving loss of life or serious life-threatening injuries.

TABLE 6: DEFINING RISK LEVELS - SEVERITY CONTINUED

<i>Qualitative Values</i>	<i>Qualitative Values</i>		<i>Description</i>
Moderate	26-75	5	The failure could be expected to have a serious adverse effect on organizational operations, organizational assets, individuals, other organizations, or the Nation. A serious adverse effect means that, for example, the failure event might: (1) cause a significant degradation in mission capability to an extent and duration that the organization is able to perform its primary functions, but the effectiveness of the functions is significantly reduced; (2) result in significant damage to organizational assets; (3) result in significant financial loss; or (4) result in significant harm to individuals that does not involve loss of life or serious life-threatening injuries.
Low	6-25	2	The failure could be expected to have a limited adverse effect on organizational operations, organizational assets, individuals, other organizations, or the Nation. A limited adverse effect means that, for example, the failure event might: (1) cause a degradation in mission capability to an extent and duration that the organization is able to perform its primary functions, but the effectiveness of the functions is noticeably reduced; (2) result in minor damage to organizational assets; (3) result in minor financial loss; or (4) result in minor harm to individuals.
Very Low	0-5	0	The failure could be expected to have a negligible adverse effect on organizational operations, organizational assets, individuals, other organizations, or the Nation.

TABLE 7: MATRIX FOR ASSESSING RISK LEVELS: LIKELIHOOD V. SEVERITY SCALE

<i>Likelihood of Impact</i>	<i>Severity of Impact</i>				
	<i>Very Low</i>	<i>Low</i>	<i>Moderate</i>	<i>High</i>	<i>Very High</i>
<i>Very High</i>	Low	Low	Moderate	High	Very High
<i>High</i>	Low	Low	Moderate	High	Very High
<i>Moderate</i>	Very Low	Low	Moderate	Moderate	High
<i>Low</i>	Very Low	Low	Low	Moderate	Moderate
<i>Very Low</i>	Very Low	Very Low	Very Low	Low	Low

APPENDIX D:

BALANCING ACTION VERSUS INACTION

After determining adjusted risk levels, decision-makers should balance the risk associated with action against the risks associated with inaction.

Consider the purpose and necessity of the technology – what are the potential benefits of the technology? What is the likelihood that the technology will succeed (although note that failures are also a critical positive feature of innovation and emerging technology like AI)? What is the likelihood that the technology or failed activity will produce useful knowledge? How much potential benefit, what kind, and to whom? How much potential harm, what kind, and to whom?

Action has risks; inaction also has risks. What is the potential harm of choosing to do nothing? Who bears that harm?

There are many potential risks of inaction that may need to be considered in the context of a specific technology. Some include:

- The inability to identify new potential connections and linkages across data and targets at the speed of operations.
- The inability to quickly identify a change in an adversary's nuclear posture due to the adversary using new techniques or deceptive practices not yet understood by human analysts.
- The inability to quickly recognize a threat or implication based on patterns associated with a target in unrelated activities.
- The inability to quickly identify voice cuts from a high value target that are not in the target's native language.
- The inability to quickly identify a potential threat at an IC facility based on slightly modified patterns of activity or suspected associates.

While these and other risks of inaction could have serious implications for the ultimate intelligence and insights the IC can provide to policymakers, these risks will not outweigh the risk of moving forward with a technology in all circumstances. In addition, the risks of inaction might be mitigated if there are alternative means of getting equivalent or similar information.

APPENDIX E: EXAMPLES

Scenario 1:

IC Agency X conducts video surveillance of its own property for force protection and security purposes. Agency X has buildings and physical infrastructure globally. Agency X has the opportunity to incorporate a new AI algorithm that conducts facial recognition into its surveillance system. The technology uses facial recognition algorithms to connect people walking or driving near the Agency's compounds to previously collected intelligence data, including people in the backgrounds of photos. The AI algorithm does not distinguish between US and non-US persons.

Agency X's Director has asked to better understand the risks associated with this technology before going forward with it. The Agency's Security Director uses the IC AI risk framework to provide additional information that will inform the Director's decision.

Step 1:

Identify potential failures and range of effects for specific AI tool or technology.

Possible First Order Risks:

- Supply Chain Security: Software supply chain includes Israeli company.
- Operational, Data: Potential for human error in storing and using the data properly without specific, newly created training.
- Intentional Adversarial Attack: Israeli company could intentionally target capability to infiltrate organization.
- Data: Potential collection of significant amounts of US persons (USP) data, especially in domestic locations. Potential for incorrect storage, access, retention, dissemination, or disposition of that data.

Possible Second Order Risks:

- Compliance, Privacy & Civil Liberties: If data is used or stored improperly, there could be compliance and P&CL implications for USP data.
- Workforce: The new training required could take time for the workforce to fully understand and implement, leading to errors.
- Reputation & Trust: If USP data is improperly handled, it would significantly damage the organization's reputation and the public's trust of the organization.
- Legal: If USP data is improperly handled, the organization could lose its authority to carry out these security monitoring activities.

Step 2:

Analyze failures and identify risks by assessing organization's vulnerabilities and range of potential effects.

- Significant amounts and different types of data are aggregated within Agency X's IT systems. If there is a supply chain security breach, the malware could infect many of the Agency X's systems, requiring significant time and resources to clean and re-secure the Agency X's IT systems and severely affecting whatever operations are connected to those systems.
- Agency X's domestic facilities are highly likely to have USPs around the perimeter. Without proper training and safeguards, USP information could be improperly handled. This could lead to significant concern from oversight entities and the public, as well as negatively impact Agency X's reputation and authority to use this AI algorithm. Even with training, human error could lead to same result.
- Operational error or adversarial attack could lead to individuals being incorrectly identified and connected to existing intelligence. An adversarial attack also could allow an adversary to more easily identify US intelligence officers entering and exiting Agency X facilities.

Step 3:

Assess, prioritize, and reconcile risks by analyzing likelihood (Table 5) and severity (Table 6) of potential impacts.

Step 3A:

Calculate initial risk level by assigning risk levels to each threat and combining the likelihood of impact with the severity of impact.

<i>Risk:</i>	<i>Likelihood:</i>	<i>Severity:</i>	<i>Risk Level:</i>
Supply Chain	Low	High	Low
Data handling	Moderate	High	Moderate
Intentional Adversarial Attack	Moderate	Very High	High
Compliance, Privacy & Civil Liberties	High	High	High
Workforce	Moderate	Moderate	Moderate
Reputation & Trust	Moderate	High	Moderate
Legal	Moderate	High	Moderate

Initial Risk Level: **High**

Step 3B:

Identify and assess potential mitigations.

Possible Mitigations:

- Data Handling and Segregation: Holding any data collected by this system separately from other agency data holdings, and only allowing it to touch during the active search for related intelligence but not to reside in any other databases, could mitigate the harm that could arise should USP data be wrongfully collected.
- Supply Chain: Security testing and other modifications could be made through a U.S. partner supplier to lower security threat. If other suppliers exist, back up suppliers could be asked to support if the security environment necessitates a change in supplier.

- **Intentional Adversarial Attack:** Modifications and review through a partner supplier could reduce vulnerability to adversarial attack. Software could be run on its own secured network to reduce exposure and vulnerability to other sensitive IC IT systems.
- **Legal:** The legal and policy offices could advise on needed policy and guidance for the program to reduce the likelihood of confusion or inconsistent/incorrect implementation.
- **Workforce:** Comprehensive training programs modeled from existing compliance training for sensitive IC systems could be created and adapted for the new software and its application. The compliance office could regularly review and audit operational implementation.
- **Reputation & Trust:** Agency X could brief Congress in advance to ensure awareness and support for this activity. The ODNI could publicly announce that some agencies in the IC are implementing force protection software that includes facial recognition capabilities to help avoid surprise without divulging which agencies are doing this.

Step 3C:

Re-calculate the level of adjusted risk factoring in mitigation plan.

<i>Risk:</i>	<i>Likelihood:</i>	<i>Severity:</i>	<i>Risk Level:</i>
Supply Chain	Low	Moderate	Low
Data Handling	Low	Moderate	Low
Intentional Adversarial Attack	Low	Very High	Moderate
Compliance, Privacy & Civil Liberties	Moderate	Moderate	Moderate
Workforce	Low	Moderate	Low
Reputation & Trust	Moderate	Moderate	Moderate
Legal	Moderate	Moderate	Moderate

Adjusted Risk Level with Mitigation Steps: **Moderate**

Step 4:

Proper decision-maker decides whether to mitigate, transfer, avoid, or accept the risk.

Possible Scenario: Since Agency X finds the objective to be a high priority and the risks can be adequately mitigated, Agency X decides to proceed with implementing the program with the below specified mitigation plan.

Step 4A:

Balance adjusted risk level against risk of inaction and opportunity costs.

If Agency X does not move forward with this new AI algorithm, current force protection activities will remain in place and continue providing good security. However, no security measures are 100% effective and this new algorithm would add an additional layer of protection against bad actors slipping through perimeter security to try to enter Agency X facilities. It would also alert security guards to the presence of someone questionable before that person approached, giving guards more time to prepare for the encounter. Without the algorithm, bad actors might be able to approach initial guarded checkpoints without any warning to those

guards, adding risk to those encounters and others at the outer edges of the facilities.

Step 4B:

Create an action plan.

- The CIO creates a data segregation plan.
- The legal and policy offices outline required policy and guidelines.
- The security office begins reviewing alternative suppliers, creates a training program (with the help of operators and lawyers), and creates a congressional briefing plan (with the legislative affairs office).
- The public affairs office considers, along with the security office and the Director of Agency X, whether to ask the DNI to make a public announcement.

Step 4C:

Communicate and share risk assessment results.

- Agency X circulates risk assessment and proposed mitigations to key agency and oversight stakeholders.
- Agency X executives share program details with other IC agencies that may be implementing similar programs to ensure awareness, consistency, and lessons learned.

Step 5:

Maintain, monitor, and periodically update the risk assessment; Evaluate results on ongoing basis.

- The compliance office creates program evaluation and assessment schedules.
- Program Managers reconvene regularly to implement findings from the security and compliance offices.
- Director of Security updates risk assessment annually.

Scenario 2:

IC Agency Y's mission is to identify, collect, and analyze foreign terrorism information. Agency Y has an opportunity to incorporate a new AI algorithm that reviews current Agency Y data holdings and learns from connections previously made by counterterrorism analysts, to identify new connections as new data is acquired.

Agency Y's Director has asked to better understand the risks associated with this technology before going forward with it. Agency Y's Counterterrorism Director uses the IC AI risk framework to provide additional information that will inform the Director's decision.

Step 1:

Identify potential failures and range of effects for specific AI tool or technology.

Possible First Order Risks:

- **Partner, Security:** The AI software supplier is a new startup that has delivered impressive initial results but lacks deep experience. There is also a very small amount of Chinese investment in the startup.
- **Intentional Adversarial Attack:** Chinese investors could try to take action to deny, disrupt, or degrade the capability in the U.S. government's application.
- **Technology:** Due to supplier's inexperience, operational performance may not be as expected.
- **Funding, Schedule:** Newness of software could result in schedule delays and cost increases.
- **Data:** Similarities in original data connections could introduce incorrect bias into the algorithm (i.e., if many of the original connections have Arabic names, the algorithm could look for only Arabic names).
- **Policy:** There are unclear or insufficient standards and policies in place at Agency Y on the use and application of this type of AI application.

Possible Second Order Risks:

- **Transparency, Work Product & Explainability:** Early demonstrations show there is difficulty in understanding and tracing the AI methodology and algorithm to identify and understand the program's decision making. Output could be biased or incorrect.

Step 2:

Analyze failures and identify risks by assessing organization's vulnerabilities and range of potential effects.

- Chinese investors could attempt to leverage their investment to influence the AI partner and gain access to sensitive information about the AI software capabilities, allowing China to glean insights on how to manipulate the algorithm's methodology to deny, disrupt, or degrade the capability, leading to inaccurate results or inoperable software.
- Technology may not ultimately work due to AI partner's inexperience, causing a decrease in Agency Y's operational readiness and capacity and a loss of taxpayer dollars.
- Undetected bias in algorithm could cause a severe impact on the quality and credibility of Agency Y's analysis and findings.

- Decision makers who rely on Agency Y's faulty analysis could make decisions that adversely impact U.S. goals and priorities.

Step 3:

Assess, prioritize, and reconcile risks by analyzing likelihood (Table 5) and severity (Table 6) of potential impacts.

Step 3A:

Calculate initial risk level by assigning risk levels to each threat and combining the likelihood of impact with the severity of impact.

<i>Risk:</i>	<i>Likelihood:</i>	<i>Severity:</i>	<i>Risk Level:</i>
Partner	Low	High	Low
Intentional Adversarial Attack	Moderate	Very High	Moderate
Technology	Moderate	Very High	High
Funding, Schedule	Moderate	Moderate	Moderate
Data	Moderate	Very High	High
Policy	Moderate	Moderate	Moderate
Transparency, Work Product & Explainability	Very High	Very High	Very High

Initial Risk Level: **High**

Step 3B:

Identify and assess potential mitigations.

Possible Mitigations:

- Testing and Evaluation: The supplier and the agency could do significant testing, evaluation, validation, and verification (TEVV) of the technology and the algorithms on synthetic data to ensure it works as expected. The TEVV process may also help the supplier develop additional mechanisms for transparency and explainability, as well as counter potential bias.
- Legal: The legal and policy offices could advise on creating new policy and guidelines that would provide sufficient standards and guidance on the use and application of this AI algorithm.
- Partner: Alternative suppliers could be recruited as backups, if any exist.
- Security: Supplier control mechanisms can be devised and implemented to reduce risks posed by Chinese investment and exploitation.
- Technology: The supplier and the agency could agree to specific gates and milestones in the development of the technology, without committing fully until those gates and milestones are met.

Step 3C:

Re-calculate the level of adjusted risk factoring in mitigations.

<i>Risk:</i>	<i>Likelihood:</i>	<i>Severity:</i>	<i>Risk Level:</i>
Partner	Low	Low	Low
Intentional Adversarial Attack	Moderate	High	Moderate
Technology	Moderate	High	Moderate
Funding, Schedule	Low	Moderate	Low
Data	Moderate	High	Moderate
Policy	Moderate	Moderate	Moderate
Transparency, Work Product & Explainability	Very High	Very High	Very High

Adjusted Risk Level with Mitigation Steps: **High**

Step 4:

Proper decision-maker decides whether to mitigate, transfer, avoid, or accept the risk.

Possible Scenario: Agency Y weighs the program's risk level and mitigation plans and alternative options. After considering the mitigation plans and reduced risk level, Agency Y decides the risk level remains too high, chooses not to proceed with implementing the program, and continues to seek an alternative option with less risk.

Step 4A:

Balance adjusted risk level against risk of inaction and opportunity costs.

Without this new algorithm, Agency Y's analysts will not be able to more quickly identify new terrorist connections. Agency Y analysts will have to continue to use older methods for identifying those connections, which in addition to being slower, also could be faulty as they are subject to human error. However, the potential for inaccurate connections being made with this new algorithm and the inability to understand or trace the algorithm's process make this software too risky to use.

Step 4B:

Create an action plan.

N/A in this scenario.

Step 4C:

Communicate and share risk assessment results.

Engage with customers and Agency Y Director to communicate assessment results and decision.

Step 5:

Maintain, monitor, and periodically update the risk assessment; Evaluate results on ongoing basis.

N/A in this scenario.

Scenario 3:

IC Agency Z's mission is to identify secret nuclear facilities in Iran and North Korea. Agency Z is presented with an opportunity to incorporate a new AI capability that uses a visual tracking algorithm to recognize existing, and learn new, patterns and indicators related to the building and/or operating of secret nuclear facilities, and alert users to those indicators.

Agency Z's Director of Analysis has asked to better understand the risks associated with this technology before going forward with it. The Agency's Director of Nonproliferation uses the IC AI risk framework to provide additional information that will inform the Director of Analysis's decision.

Step 1:

Identify potential failures and range of effects for specific AI tool or technology.

Possible First Order Risks:

- **Data & Algorithms:** Data on this issue is scarce and connections are tenuous resulting in potential incorrect assumptions by the algorithm.
- **Technology:** Requirements have shifted significantly during development.
- **Schedule, Funding:** Changing requirements have resulted in slight delay and increased costs.

Possible Second Order Risks:

- **Transparency:** Analysts cannot trace methodology of algorithms.
- **Work Product & Explainability:** Output could be incorrect; workforce may be unable to explain or interpret output.

Step 2:

Analyze failures and identify risks by assessing organization's vulnerabilities and range of potential effects.

- Inaccurate results from use of the AI algorithm could reduce the credibility of Agency Z's analysis and its mission to provide decision makers with critical information.
- Decision makers who rely on Agency Z's faulty analysis could make decisions that adversely impact U.S. goals and priorities.
- Shifting requirements during development could cause the algorithm to perform inconsistently through development, rendering the use of the algorithm unreliable to Agency Z.

Step 3:

Assess, prioritize, and reconcile risks by analyzing likelihood (Table 5) and severity (Table 6) of potential impacts.

Step 3A:

Calculate initial risk level by assigning risk levels to each threat and combining the likelihood of impact with the severity of impact.

<i>Risk:</i>	<i>Likelihood:</i>	<i>Severity:</i>	<i>Risk Level:</i>
Data & Algorithms	Moderate	Very High	High
Technology	Moderate	High	Moderate
Schedule	High	Low	Low
Transparency	High	Very High	Very High
Work Product & Explainability	Moderate	Very High	High

Initial Risk Level: **High**

Step 3B:

Identify and assess potential mitigations

Possible Mitigations:

- Schedule & Technology: Program Managers, with operators and technologists, could set baseline requirements for an AI solution that has the capacity to be grown to a tailored focus but starts development as an adaptable platform.
- Testing and Evaluation: The supplier and the agency could do significant testing, evaluation, validation, and verification (TEVV) of the technology and the algorithms on past confirmed nuclear-related activities to ensure technology works as expected.
- Transparency: The AI supplier could enter into an iterative, collaborative process with analysts to understand algorithm requirements, priorities, and operations. The TEVV process may also help the supplier develop additional mechanisms for transparency and explainability.
- Work Product & Explainability and Data & Algorithms: Analysts could regularly audit results of technology with older, trusted analytic techniques until appropriate levels of confidence in technology are reached.

Step 3C: Re-calculate the level of adjusted risk factoring in mitigations

<i>Risk:</i>	<i>Likelihood:</i>	<i>Severity:</i>	<i>Risk Level:</i>
Data & Algorithms	Moderate	High	Moderate
Technology	Low	High	Low
Schedule	Low	Low	Low
Transparency	High	High	High
Work Product & Explainability	Low	High	Low

Adjusted Risk Level with Mitigation Steps: **Moderate**

Step 4:

Proper decision-maker decides whether to mitigate, transfer, avoid, or accept the risk.

Possible Scenario: Due to the high priority of this objective, the lack of alternative options, and the ability to adequately mitigate the identified risks, Agency Z proceeds with implementing the program with mitigations.

Step 4A:

Balance adjusted risk level against risk of inaction and opportunity costs.

Without this new capability, analysts will have to continue to use old methods which are slow, difficult, and imperfect. There will be continued risk of missed indicators and undetected nuclear activities that could result in policymaker surprise. This capability presents a very big step forward in better identifying and tracking illicit nuclear programs.

Step 4B:

Create an action plan.

- TEVV process includes routinely testing algorithm against past known nuclear activities of interest until confidence is reached.
- Program Managers, with operators and technologists, set baseline requirements for capability and plan for working on transparency with supplier.
- Analysts create audit process to ensure algorithm working as expected.

Step 4C:

Communicate and share risk assessment results.

- Director of Nonproliferation communicates expectations and risks to leadership and approves baseline requirements.
- Agency Z executives share program details with other stakeholders, as needed.

Step 5:

Maintain, monitor, and periodically update the risk assessment; Evaluate results on ongoing basis.

- The compliance office creates program evaluation and assessment schedules.
- Analysts implement audit process to confirm accuracy of data output and methodology.
- Director of Nonproliferation updates risk assessment annually.



TECH, LAW & SECURITY
PROGRAM