



# PROCEEDINGS OF THE TWENTY-THIRD ANNUAL ACQUISITION RESEARCH SYMPOSIUM AND INNOVATION SUMMIT

---

WEDNESDAY, MAY 6, 2026 SESSIONS  
VOLUME I  
“ACCELERATING WARFIGHTING CAPABILITIES”

Approved for public release; distribution is unlimited.

Prepared for the Naval Postgraduate School, Monterey, CA 93943.

Disclaimer: The views expressed are those of the author(s) and do not reflect the official policy or position of the Naval Postgraduate School, US Navy, Department of Defense, or the US government.



# Preface & Acknowledgements

---

The purpose of the Accelerating Warfighting Capabilities, NPS 23rd Annual Acquisition Research Symposium and Innovation Summit is to provide a forum for the presentation of scholarly acquisition research, as well as for dialogue between scholars and acquisition policy-makers and practitioners. Research papers and presentations are given on recently completed and on-going Departments of Defense and US Navy (DoD/DON)-sponsored projects conducted by researchers at a variety of research institutions. Senior DoD/DON acquisition officials serve as panelists or keynote speakers to present their critiques and comments on research papers and priorities.

This year our symposium is coupled with an Innovation Summit and takes up the theme of “Accelerating Warfighting Capabilities.”

Although attendees come from many U.S. locations, as well as from some international locales, a large number are from Naval Postgraduate School (NPS) where faculty members and graduate students engage in acquisition-related research. In particular, NPS graduate students are an integral component of the research and dialogue. The Symposium serves an essential part of their graduate learning experience and provides them the opportunity to meet with senior policy-makers, practitioners, and distinguished scholars.

We gratefully acknowledge the ongoing support and leadership of our sponsors, whose foresight and vision have assured the continuing success of the Acquisition Research Program:

- Assistant Secretary of the Navy for Research, Development and Acquisition (ASN (RDA))
- Director, Acquisition Talent Management, U.S. Navy (DATM)
- Program Executive Officer, Integrated Warfare Systems (PEO IWS)

The research presented in this report was supported by the Acquisition Research Program, Department of Acquisition, Finance, and Manpower at the Naval Postgraduate School.

**To request defense acquisition research, please contact:**

Acquisition Research Program  
Department of Acquisition, Finance, and Manpower  
Naval Postgraduate School  
E: [arp@nps.edu](mailto:arp@nps.edu)  
[www.acquisitionresearch.net](http://www.acquisitionresearch.net)

Copies of Symposium Proceedings and Presentations, and Acquisition Sponsored Faculty Reports, and Student Research Reports and Posters may be viewed and printed from the **NPS Defense Acquisition & Innovation Repository** at <https://dair.nps.edu/>



ACQUISITION RESEARCH PROGRAM  
DEPARTMENT OF ACQUISITION, FINANCE, AND MANPOWER  
NAVAL POSTGRADUATE SCHOOL

# Table of Contents

---

|                                                                                                                                                                       |            |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| <b>WELCOME: ANN E. RONDEAU, ED.D, VADM, U.S. NAVY (RET.), PRESIDENT, NAVAL POSTGRADUATE SCHOOL.....</b>                                                               | <b>1</b>   |
| <b>WELCOME: JASON L. POTTER, PERFORMING THE DUTIES OF THE ASSISTANT SECRETARY OF THE NAVY (RESEARCH, DEVELOPMENT AND ACQUISITION) .....</b>                           | <b>3</b>   |
| <b>PANEL 1. ACCELERATING WARFIGHTING CAPABILITIES: PUTTING THE INDUSTRIAL BASE ON A WARTIME FOOTING .....</b>                                                         | <b>5</b>   |
| <b>PANEL 2. ENGINEERING THE SOFTWARE-DEFINED ENTERPRISE: AGILE PATHWAYS, MBSE, AND AI ASSURANCE .....</b>                                                             | <b>9</b>   |
| The Pentagon's Revolution in Software-Defined Warfare and Its Testing Dilemma .....                                                                                   | 11         |
| Digital Engineering to Improve the Quality of a Capability Needs Statement .....                                                                                      | 34         |
| Transforming Acquisition for Speed, Agility, and Effectiveness .....                                                                                                  | 52         |
| <b>PANEL 3. ADAPTIVE LETHALITY: SYNCHRONIZING DIGITAL TWINS AND EMBEDDED ENGINEERING FOR RAPID FIELDING .....</b>                                                     | <b>65</b>  |
| Rapid Deployment of Airborne Directed Energy Weapons: A Strategic Solution to Evolving Aerial Threats .....                                                           | 67         |
| Embedded Capability Development: A Case Study in Rapid Missile Prototyping and Transition..                                                                           | 85         |
| Cyber Digital Twin-Informed Zero Trust: A Synergistic Framework for Securing Operational Technology in Defense Logistics Infrastructure.....                          | 103        |
| <b>PANEL 4. THE FRICTIONLESS FRONTIER: REFORMING ACQUISITION TO SECURE CUTTING-EDGE CAPABILITY .....</b>                                                              | <b>119</b> |
| Exploring the Use Case of Wargaming Tools Across the Contracting Life Cycle .....                                                                                     | 120        |
| Is the Military Attracting New Companies? Assessing First-Time Entrants into the Defense Market                                                                       | 134        |
| Buying Flexibility: Converting Cost-Plus to Firm-Fixed-Price to Strengthen Mission Resilience .                                                                       | 149        |
| <b>PANEL 5. HIGH-GROUND INNOVATION: SYNCHRONIZING AGILE DEVELOPMENT AND PRIVATE INVESTMENT IN SPACE.....</b>                                                          | <b>163</b> |
| Accelerating DoW Space-Based Acquisition Through the Implementation of Agile, DevSecOps and Digital Engineering Processes .....                                       | 165        |
| Aligning Defense Products to National Security Space Needs .....                                                                                                      | 179        |
| <b>PANEL 6. GUARDING THE MACHINE: PRESERVING HUMAN AGENCY AND MARKET INTEGRITY IN THE AI ERA .....</b>                                                                | <b>193</b> |
| Keeping Humans in Command: Preserving Human Agency in AI-Driven Federal Acquisition ....                                                                              | 194        |
| AI Collusion in Procurement Auctions .....                                                                                                                            | 218        |
| An Actionable Framework & Interactive Web Application for Transitioning Technology to the Warfighter: The Transition Maturity Framework (TMaF) and Tech2PEO/PAE ..... | 227        |
| <b>PANEL 7. THE ARSENAL OF INNOVATION: REVITALIZING INDUSTRIAL SURGE AND MANUFACTURING SECURITY .....</b>                                                             | <b>231</b> |
| Connecting the Kill Chain to the Supply Chain: Building Industrial Surge Capacity .....                                                                               | 232        |



|                                                                                                                                                                                             |            |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| Establishing a Manufacturing Security Program.....                                                                                                                                          | 249        |
| The Gears of War: The State of the Aerospace Gear Industry and the Implications for Accelerating Acquisition.....                                                                           | 269        |
| <b>PANEL 8. WINNING THE COMPETITION: LEVERAGING ACQUISITION FOR STRATEGIC ADVANTAGE .....</b>                                                                                               | <b>283</b> |
| From Policy to Practice: Drivers of Acquisition Reform .....                                                                                                                                | 284        |
| “东风-17”: 高超声速让反导系统形同虚设 [“Dongfeng-17:” Hypersonics Render Missile Defense Systems Useless].” (January 15, 2018). Science and Technology Daily. ....                                         | 310        |
| Defense Acquisition and Contracting Approaches: Implications for Foreign Military Sales .....                                                                                               | 311        |
| <b>PANEL 9. FUTURE-PROOFING THE LIFECYCLE: SILO-BUSTING THROUGH INTEGRATED DIGITAL MATURITY .....</b>                                                                                       | <b>325</b> |
| Quantum Acquisition: A New Paradigm for Understanding the Interdependencies of Complex Networks in the Acquisition Life Cycle of Warfighting Systems.....                                   | 326        |
| AI-Infused Integrated Digital Maturity Pathway (IDMP) User-Story Framework for Defense Acquisition Transformation.....                                                                      | 339        |
| Augmented Acquisition: Operationalizing Human-in-the-Loop AI to Accelerate Warfighting Capabilities .....                                                                                   | 356        |
| <b>PANEL 10. EMPOWERING THE PROFESSIONAL: SHAPING THE FUTURE OF ACQUISITION TALENT .....</b>                                                                                                | <b>371</b> |
| <b>PANEL 11. OPPORTUNITIES AND CONSIDERATIONS FOR ACQUISITION IN THE SECURITY COOPERATION SPACE .....</b>                                                                                   | <b>375</b> |
| Improving the DoW’s Ability to Leverage Allied and Partner Capabilities for Burden Sharing .....                                                                                            | 376        |
| Total Demand: Integrating Foreign Military Sales and Domestic Acquisition to Fortify the Defense Industrial Base and Strengthen Alliances .....                                             | 388        |
| Embracing Non-Traditional Solutions: Evolving Security Cooperation for the Modern Threat Environment.....                                                                                   | 401        |
| <b>PANEL 12. BEYOND THE LINEAR BARRIER: DATA-DRIVEN STRATEGIES FOR ACCELERATED ACQUISITION.....</b>                                                                                         | <b>411</b> |
| A Study on Establishing a Rapid Acquisition Process for the Military Application of Advanced Commercial Technologies in Korea .....                                                         | 413        |
| Recent Improvements in Acquisition Speed to Initial Operational Capability .....                                                                                                            | 428        |
| Weapon Systems Annual Assessment: DoD Leaders Should Ensure That Newer Programs Are Structured for Speed and Innovation .....                                                               | 438        |
| <b>PANEL 13. FROM DEMO TO DETERRENCE: ACCELERATING SURGE CAPACITY AND FLEET INTEGRATION .....</b>                                                                                           | <b>441</b> |
| Pathfinding Innovation for Legacy Munitions .....                                                                                                                                           | 442        |
| Generating Adaptive Capability from Constraints: Aligning Operational Dependency, Requirements, and Competitive Procurement to Sustain Advantage in Maritime Expeditionary Operations ..... | 462        |
| Demo to Deployment in 30 Days: The ARMS Case Study.....                                                                                                                                     | 482        |





## **WELCOME: ANN E. RONDEAU, ED.D, VADM, U.S. NAVY (RET.), PRESIDENT, NAVAL POSTGRADUATE SCHOOL**

---



**Ann E. Rondeau, Ed.D, Vice Admiral, U.S. Navy (Ret.),** was appointed as President, Naval Postgraduate School on January 29, 2019. She brings to the assignment an unparalleled record of leadership and achievement within the military and academia in the areas of education, training, research, executive development, change management, and strategic planning. Prior to her appointment, Adm. Rondeau served as the sixth president of the College of DuPage. Her most recent military position was as the President of the National Defense University, a consortium of five colleges and nine research centers in Washington, DC.

Rondeau has extensive leadership experience in significant military and educational roles. In 1985, she was selected and served as a White House Fellow in the Reagan Administration and went on to serve as the Deputy Commander of the U.S. Transportation Command in Illinois, Pentagon Director/Chief of Staff for the U.S. Navy Staff, Commander of the Navy Personnel Development Command in Virginia, Commander of the Naval Service Training Command at Great Lakes, Ill., Pacific Fleet Staff Chief of Staff in Hawaii, Commanding Officer of Naval Support Activity in Tennessee and other staff and commanding responsibilities with policy, planning, Fleet support, joint logistics, training and education. Rondeau retired from the U.S. Navy as a three-star admiral in 2012 and was the second woman to have achieved that rank in the Navy. She then served as a partner and later an independent consultant with the IBM Watson group.

President Rondeau's leadership has served many, both past and present, to include: Board of Directors, United States Institute of Peace; Board of Directors, German Marshall Fund; Board of Directors, The Atlantic Council; Board of Directors, National Museum of the American Sailors; Board of Directors, Council of Higher Education Accreditation; Board of Directors, Chicago Regional Growth Corporation; Board of Directors, Choose DuPage (regional development organization for Chicago northwest suburbs); Tennessee/Mid-South Economic Development Board; DoD liaison to the Center for the Study of the Presidency; Military Advisory Board (studying energy and environment impacts on national security); Flag Officer Advisory Council for Arizona State University, the National Naval Officers Association Senior Advisory Panel, the Eisenhower Memorial Commission and the National Cold War Veterans Memorial Design Steering Committee among others.

Rondeau holds a B.A. from Eisenhower College (NY), an M.A. from Georgetown University (DC) and an Ed.D. from the College of Education at Northern Illinois University in DeKalb. She also holds an honorary Doctorate in Public Service from Carthage College (Kenosha, WI) and an honorary Doctorate in Humane Letters from Rosalind Franklin University of Medicine and Science (Chicago, IL).



THIS PAGE INTENTIONALLY LEFT BLANK



## **WELCOME: JASON L. POTTER, PERFORMING THE DUTIES OF THE ASSISTANT SECRETARY OF THE NAVY (RESEARCH, DEVELOPMENT AND ACQUISITION)**

---



**Jason L. Potter, Performing the duties of the Assistant Secretary of the Navy (Research, Development and Acquisition):** In July 2025, Mr. Jason Potter was assigned to perform the duties of the Assistant Secretary of the Navy (Research, Development and Acquisition). In this role, he leads the Department of the Navy's (DON) Research, Development, Acquisition, and Sustainment programs and the DON's contracting community.

Mr. Jason Potter was appointed Deputy Assistant Secretary of the Navy for Ship Programs in May 2025. He is also performing the duties of the Principal Civilian Deputy Assistant Secretary of the Navy (Research, Development and Acquisition).

Prior to his appointment, Mr. Potter served as a professional staff member in the United States Senate for 10 years on the Appropriations and Armed Services Committees. His responsibilities included oversight of Navy shipbuilding; Navy and Air Force procurement accounts; various research, development, test, and evaluation accounts; Defense Production Act purchases; National Defense Stockpile transactions; and defense acquisition law.

Earlier in his career, Mr. Potter served in civilian positions on the Navy Staff (OPNAV) including as a special assistant to the Chief of Naval Operations, programming division section head, and force structure analyst. He also served as a naval officer with assignments as Navigator aboard USS CAPE ST GEORGE (CG 71), Communications and Repair Officer aboard USS RUSHMORE (LSD 47), naval liaison to Joint Special Operations Task Force – Philippines, and analyst on the 2010 Navy Quadrennial Defense Review Integration Group. He is currently serving as a Commander in the U.S. Navy Reserve.

Mr. Potter holds an executive M.B.A. from the Naval Postgraduate School, a M.A. in Government from the University of Maryland, and a B.S. in Political Science from the U.S. Naval Academy.



THIS PAGE INTENTIONALLY LEFT BLANK



## PANEL 1. ACCELERATING WARFIGHTING CAPABILITIES: PUTTING THE INDUSTRIAL BASE ON A WARTIME FOOTING

Wednesday, May 6, 2026, 1015 – 1130 ET (0715 - 0830 PT)

### Panel Summary:

As the global security environment becomes increasingly dangerous, the United States can no longer rely on existing stockpiles or ossified acquisition patterns to prevail in industrial-scale warfare. This plenary panel, serving as the symposium's kickoff, explores the department-wide transformation required to shift the entire acquisition system onto a "wartime footing". By examining the intersection of industrial policy, surge capacity, and Indo-Pacific security, these top Pentagon leaders will discuss concrete actions being taken to rebuild the arsenal of freedom and deliver capabilities at the speed and scale required for modern conflict.

**Chair: Dr. Jerry McGinn**, Director, Center for the Industrial Base and Senior Fellow, Defense and Security, Center for Strategic and International Studies (CSIS)

### Panelists:

**Synchronizing Global Logistics and Sustainment** – *VADM Dion English, Director of Logistics at the Joint Staff*

**Strategic Industrial Base Policy and Surge Capacity** – *Honorable Michael Cadenazzi, Assistant Secretary of War for Industrial Base Policy*

**Modernizing the Acquisition Ecosystem** – *Mr. James Ruocco, Senior Executive Service, performing duties of Assistant Secretary of War for Acquisition*



**Dr. Jerry McGinn**—is the director of the Center for the Industrial Base and a senior fellow with the Defense and Security Department. He is one of the nation's leading experts on industrial base, acquisition, government contracting, and security policy issues. Prior to joining CSIS, Jerry served as executive director of the Greg and Camille Baroni Center for Government Contracting at George Mason University, publishing over 70 influential reports, white papers, and commentaries. Prior to GMU, Jerry served as the senior career official in the Office of Manufacturing and Industrial Base Policy in the Department of Defense (DoD), leading efforts to analyze the capabilities of the defense industrial base. He also directed hundreds of reviews of high-profile mergers and acquisitions as well as transactions before the Committee on Foreign Investment in the United

States. Before his service in DoD, Dr. McGinn spent a decade in senior defense industry roles at Northrop Grumman, Deloitte Consulting LLP, and QinetiQ North America. Prior to industry, Dr. McGinn served as special assistant to the principal deputy undersecretary for policy in DoD and as a political scientist at RAND. Dr. McGinn is also a widely acclaimed thought leader and has testified numerous times before the U.S. Congress and the UK House of Commons. Dr. McGinn served with distinction as an infantry officer and is a graduate of Ranger and Airborne Schools. He earned a PhD, MS, and MA from Georgetown University as well as a BS from the United States Military Academy.





**Vice Admiral Dion D. English, USN**—is a native of Monroeville, Alabama, and of Slidell, Louisiana. English entered the Navy's Broadened Opportunity for Officer Selection and Training program in 1987. In 1993, he graduated from Louisiana State University and was subsequently commissioned through the Navy Reserve Officers' Training Corps program at Southern University. He earned master's degrees in Business Administration from Old Dominion University and in National Security and Strategic Studies from the Naval War College. He is also a 2007 graduate of the Joint Forces Staff College and attended the Executive Program in Strategy and Organization at the Stanford Graduate School of Business.

His afloat assignments include disbursing and sales officer, USS Scott (DDG 995); cargo and stock control officer, USNS San Diego (T-AFS 6); supply officer, USS Barry (DDG 52); and commanding officer, Explosive Ordnance Disposal Expeditionary Support Unit ONE.

His shore assignments include ship force support officer, Supervisor of Shipbuilding, Newport News, Virginia; executive assistant to the Commander, Defense Supply Center Columbus, Ohio; deputy logistics operations officer, Joint Task Force Civil Support (JTF-CS), Fort Monroe, Virginia; logistics readiness officer, Logistics Current Operations Division, N4 Directorate for Commander, U.S. Pacific Fleet (USPACFLT); director, Supply Corps Detailing, Navy Personnel Command (PERS-4412), Millington, Tennessee; commanding officer, Naval Supply Systems Command (NAVSUP) Fleet Logistics Center (FLC) Sigonella, Italy; and deputy director, supply and distribution, J44, Joint Staff, Washington, District of Columbia.

His flag assignments include director, logistics, fleet supply and ordnance, N4, USPACFLT, Pearl Harbor, Hawaii; vice director, logistics, J-4, Joint Staff; and director, logistics division, N4L, Office of the Chief of Naval Operations.

English is designated as a Surface Warfare Supply Corps Officer and as a Navy Expeditionary Supply Corps Officer. He is a member of the Defense Acquisition Corps and is a Joint Qualified Officer. While under his command, NAVSUP FLC Sigonella earned recognition in partnership with the team that received the prestigious Admiral Stan Arthur Award for 2018 Logistics Team of the Year.

His personal decorations include Defense Superior Service Medal, Legion of Merit, Defense Meritorious Service Medal, Meritorious Service Medal, Joint Service Commendation Medal, Navy Commendation Medal, Joint Service Achievement Medal, and Navy Achievement Medal.



**Honorable Michael Cadenazzi**—was sworn in as the Assistant Secretary of War for Industrial Base Policy (ASD(IBP)) on September 23, 2025. In this role, he is the principal advisor to the Under Secretary of War for Acquisition and Sustainment on industrial base policies and leads the Department of War's efforts to develop and maintain the U.S. defense industrial base to ensure a secure supply of materials critical to national security.

Over the past two decades, Mr. Cadenazzi has served as a serial entrepreneur and consultant with experience in managing and addressing challenging issues across the aerospace & defense sector. He has launched multiple defense industry services and technology start-ups across the signals intelligence, program analysis, data analytics, and market assessment and strategy sectors, and executed two successful transfers of business ownership. His extensive sector experience includes work from the space to undersea domains and from aircraft and munitions to armor, weapons, ships, and services. His clients have included domestic and international firms from large prime contractors to all levels in the supply chain. His direct experience spans strategy, mergers & acquisitions, operations, supply chain and organizational transformation.

Prior to his civilian career, Mr. Cadenazzi served for ten years as an active-duty U.S. Navy cryptologic warfare officer. He completed the Cryptologic Division Office Course at Corry Station in Pensacola Florida before his first tour at Naval Communications & Telecommunications Area Master Station (NCTAMS) WESTPAC in Agana, Guam. Following graduate education, he was assigned to the staff of Commander,



U.S. Navy SIXTH Fleet in Gaeta Italy. He completed his military service on the staff of U.S. Naval Forces Europe in London, the United Kingdom.

Mr. Cadenazzi holds a bachelor's degree in engineering from Tulane University and a master's degree in electrical engineering with an emphasis on RF communications and signals intelligence from the U.S. Naval Postgraduate School. He was commissioned an ensign in 1995 through the Tulane Naval Reserve Officer Corps (NROTC) program.



**James A. Ruocco**—is a member of the Senior Executive Service, is currently performing the duties of the Assistant Secretary of War for Acquisition (ASW(A)). In this position, he advises the Under Secretary of War for Acquisition and Sustainment (USW(A&S)), the Deputy Secretary of War, and the Secretary of War on matters relating to the Department of War (DoW) Acquisition System, acquisition program management, and the development of strategic, space, intelligence, tactical warfare, command and control, and business systems.

Most recently, Mr. Ruocco served as Deputy Assistant Secretary of Defense for Strategic, Space, and Intelligence Portfolio Management (SSIPM). He is responsible for acquisition shaping, analysis, and oversight of warfighter capability portfolios across the Department of Defense associated with nuclear weapons systems; nuclear command, control, and communications; space;

missile defense; and command, control, communications, computers, intelligence, surveillance, and reconnaissance domains.

Previously, he served as Acting Deputy Assistant Secretary for Defense for Platform and Weapon Portfolio Management and was responsible for assessing and analyzing major platforms and weapons capability portfolios across the Department, covering all cyber, electronic warfare, aircraft, ship, submarine, ground vehicle and weapons & munitions programs. Additionally, he served as Director for Air Platforms and Weapons and was responsible for portfolio and acquisition management, policy, and investments in all Department of Defense (DoD) aircraft-related missions and capabilities.

Mr. Ruocco holds a Bachelor of Science in Aerospace Engineering from Embry Riddle Aeronautical University and a master's degree in National Resource Strategy from the National Defense University. Additionally, he graduated from the Federal Executive Institute, and holds certificates from the University of Virginia and the Massachusetts Institute of Technology. He is a member of the Acquisition Professional Community with certifications in Program Management and Systems Engineering. Mr. Ruocco has received numerous Command-level recognitions and awards, as well as being a three-time recipient of the Secretary of Defense Exceptional Civilian Service Medal.





THIS PAGE INTENTIONALLY LEFT BLANK



## PANEL 2. ENGINEERING THE SOFTWARE-DEFINED ENTERPRISE: AGILE PATHWAYS, MBSE, AND AI ASSURANCE

Wednesday, May 6, 2026, 1145 – 1300 ET (0845 - 1000 PT)

### Panel Summary:

As the Department moves away from hardware-centric acquisition, software has become the primary differentiator of military preeminence. This panel explores the critical shift toward Software-Defined Warfare, demonstrating how digital engineering models and modernized acquisition pathways eliminate traditional bottlenecks. By integrating rapid adaptability with rigorous assurance, these research initiatives provide the framework necessary to deliver dependable, AI-enabled capabilities at the speed of the evolving threat.

**Chair: Thomas W. Simms**, Executive Director, Systems Engineering and Architecture, Office of the Assistant Secretary of War for Mission Capabilities

### Presenters:

**The Pentagon's Revolution in Software-Defined Warfare and Its Testing Dilemma** – *The Honorable Nickolas H. Guertin, Virginia Tech National Security Institute*

**Digital Models to Improve the Quality of Capability Needs Statements (CNS)** – *Alfred (Fred) Schenker, Carnegie Mellon University Software Engineering Institute*

**Transforming Acquisition for Speed, Agility, and Effectiveness** – *Rita Creel, Carnegie Mellon University Software Engineering Institute*



**Thomas W. Simms**—is the Executive Director for Systems Engineering and Architecture within the Department of War (DoW) Office of the Assistant Secretary of War for Mission Capabilities (OASW(MC)). He leads efforts to advance DoW practice in systems engineering, software engineering, digital engineering, and specialty engineering areas such as reliability, manufacturing, and human systems integration. He leads developmental test, evaluation, and assessment teams performing critical technical risk and engineering assessments on major acquisition programs.

Mr. Simms also leads workforce development initiatives to guide the Engineering and Technical Management (ETM) and Test and Evaluation (T&E) workforce. He oversees efforts to build the ETM and T&E training frameworks to continually improve workforce knowledge and skills.

Mr. Simms has more than 35 years of experience in engineering to include developmental test and evaluation of complex weapon systems. Before his current role, he served as the acting Principal Deputy Director for SE&A and as Deputy Director for Engineering Policy and Systems, performing executive-level duties and developing engineering policy, guidance, and processes.

Before joining the Office of the Secretary of War, Mr. Simms worked with the Missile Defense Agency (MDA), involved with the Aegis Ballistic Missile Defense program. He became the MDA Director for Test Policy and Integration, responsible for the agency's test policy and test manpower. He worked with senior leaders from across the Department on missile defense–related topics.

Mr. Simms served as the Cruise Missile Technical Advisor on the staff of the Commander, Submarine Forces Pacific, where he provided technical expertise on both the Submarine-Launched TOMAHAWK Weapon System and the Encapsulated HARPOON Weapon System. He then served as the HARPOON Weapon System engineer in the Cruise Missile Weapon Control Systems Program Office (PMA-282) and



led the engineering development and T&E of upgrades to the HARPOON Shipboard Command Launch Control Systems.

Mr. Simms earned a Bachelor of Science in Mechanical Engineering from the University of Massachusetts, Dartmouth, and a Master's Degree from the George Washington University.



# The Pentagon's Revolution in Software-Defined Warfare and Its Testing Dilemma

**Douglas C. Schmidt**—is the Dean of William & Mary's School of Computing, Data Sciences & Physics, after spending more than 20 years as a CS professor at Vanderbilt University. He recently served as the Pentagon's Director, Operational Test and Evaluation and was a program manager at DARPA and the Chief Technology Officer at Carnegie Mellon University's Software Engineering Institute. He has published many papers and books on software-related topics. He received Bachelor and Master of Arts degrees in Sociology from William & Mary and Master of Science and Doctorate degrees in Computer Science from the University of California, Irvine. [dcschmidt@wm.edu]

**Nickolas H. Guertin**—is a defense acquisition executive with more than 30 years of experience delivering complex naval, aerospace, and software-intensive systems. He served as Assistant Secretary of the Navy for Research, Development, Acquisition, and Sustainment (2023–2025), overseeing a \$130 billion portfolio and leading 130,000 professionals. Previously, as Director of Operational Test and Evaluation, he provided independent assessments to Congress and the Secretary of Defense. A former enlisted submariner and Navy Reserve Engineering Duty Officer, he is now a Virginia Tech Senior Research Fellow, advisor, and published thought leader in acquisition reform, digital engineering, and open systems. [nickolashg@vt.edu]

**John E. Robert**—is a Principal Engineer at Carnegie Mellon University's Software Engineering Institute (SEI) and Deputy Director of the Software Solutions Division, leading software engineering R&D and Navy program support teams. He co-authored *Architecting the Future of Software Engineering: A National Agenda for Software Engineering Research & Development*, a multi-year roadmap for next-generation software-reliant systems. Previously, he served as a Department of the Navy civilian in the distributed simulation team at the Naval Air Warfare Center Aircraft Division (NAWCAD). [jer@sei.cmu.edu]

## Abstract

Warfare is inherently messy and adaptive—Sun Tzu's observation that “all warfare is based on deception” remains relevant—but today's tempo of capability delivery is outpacing hardware-centric acquisition and legacy warfighting patterns. This paper argues that military preeminence increasingly depends on software-defined warfare, where code—not platforms—becomes the decisive differentiator. We characterize this shift through six tenets: rapid adaptability, AI-driven decision support, digital twins and simulation, reprogrammable weapons, autonomous systems, and cyber operations. Together, these tenets demand unprecedented operational agility, enabling forces to reconfigure tactics, platforms, and effects during conflict.

The same features that enable overmatch also introduce fragility: tightly coupled “kill webs,” vulnerabilities in AI reasoning, and the risk of cascading failure from a single software update. This creates a central Pentagon dilemma: software-enabled capabilities can be fielded faster than they can be objectively assessed. Traditional test and evaluation (T&E), optimized for static hardware designs, is straining under continuous updates and complex interdependencies. We propose a reinvention of T&E—supported by digital twins, AI-augmented testing, DevSecOps pipelines, and independent oversight—and offer recommendations to balance rapid innovation with assurance so software-defined arsenals remain agile and dependable in the fog of war.

## Introduction

At a desert airstrip, dozens of autonomous drones lift off in unison—a buzzing swarm that embodies the Pentagon's emerging operational concept and supporting acquisition reforms of software-defined warfare. Unlike past eras defined by hardware-centric superiority, today's battlespace is increasingly shaped by the code that runs in these systems. In this emerging paradigm, software upgrades, artificial intelligence (AI) models, and digital reconfigurations—in addition to new airframes, hulls, ground vehicles, and the systems they carry—dictate the pace of adaptation. U.S. Navy warships now received combat software patches mid-deployment,



uncrewed aerial systems (UAS) can rapidly be retasked with new behaviors, and AI tools can generate operational plans or prototype code in hours rather than months. These advances collectively promise unmatched agility: the ability to reprogram the rapidly evolving military capability as quickly as adversaries shift their tactics.

This revolution is needed to win future battles but also includes new challenges. In a software-defined approach, new defense system capabilities are identified based on analysis of operational data, and new capabilities are created and integrated into systems (McNamara, 2025; Mulchandani, 2022). Each tenet of software-defined warfare shown in Figure 1 also introduces enormous, and currently unavoidable, complexity into the defense ecosystem. Systems once evaluated in isolation are now interdependent nodes in sprawling kill webs. A single line of faulty code, a misaligned interface, or an unanticipated AI decision can ripple through the joint force with cascading consequences. The very features that make software-defined warfare so powerful—speed, malleability, scale—also require technology insertion at the speed of relevance (DoW, 2026) and rethinking the testing and modernization activities (DoD, 2025).

The revolution in software-defined warfare is also why the Pentagon now faces a profound testing dilemma (Schmidt & Guertin, 2025). While development and deployment cycles accelerate, the processes to verify and validate these capabilities have not kept pace. Traditional test ranges, certification checklists, and siloed evaluations were not designed for swarming drones, cyber-resilient kill webs, or continuously updated digital twins. Without rigorous and independent test and evaluation (T&E), the United States risks fielding systems that look formidable on paper but falter in the fog of war. Realizing the promise of software-defined warfare requires not only ingenuity in development but also innovation—and renewed discipline—in testing. For example, the Pentagon could increase the synergies between developmental test and evaluation (DT&E) and operational test and evaluation (OT&E) so that validated developmental test (DT) evidence can count toward operational test (OT) findings earlier in system life cycles.

This paper explores how the Pentagon is transforming its acquisition and operational models, and why ensuring trust in these software-driven systems may prove the hardest battle of all. The remainder of this paper is organized as follows: Software-Define Warfare: Key Tenets elaborates on the tenets of software-defined warfare; A Revolution in How the Pentagon Builds Technology examines AI-augmented development and accelerated acquisition practices; Fielding Faster Than We Can Test analyzes the widening gap between rapid fielding and the pace of testing; Trust, Technology, and the Testing Tightrope highlights both the enduring importance of independent comprehensive test and evaluation, along with both the promises and pitfalls of automation; Risks of Undertested Warfare describes key risks of under-tested warfare; Recommendations: Building a Test Enterprise for Software-Defined Warfare offers recommendations on how to build a resilient test enterprise for the software-defined warfare era; and Concluding Remarks presents concluding remarks.





Figure 1. Key Tenets of Software-Defined Warfare

## Software-Defined Warfare: Key Tenets

Software-defined warfare is an emerging paradigm where software rather than hardware dictates the pace, form, and outcome of modern conflict (Clark et al., 2023). This paradigm will require generalizing and expanding our thinking of military programs. We must now reorient our actions to consider the composite capability as a larger system (or more likely a system-of-systems) and how it will be built and tested over many iterations of new functionality. These products must interact across kill-webs and embrace a different design practice of fluid integration to deliver new performance. In this way, militaries can now evolve warfighting capabilities rapidly via code into targeted sub-systems (and systems). Scaling and expanding these concepts across deployment nodes for associated mission capabilities requires fundamental changes or key tenets of software-defined warfare, as shown in Figure 1 and described below:

- **Rapid adaptability**—Weapon systems and critical supporting information, like user training and updates to tactics, can be reprogrammed or updated in hours, days, or weeks rather than over decades-long development cycles. Software updates, new AI models, or modular code allow forces to swiftly adjust to emerging threats or missions. For example, a warship's critical systems can “*get smarter overnight with a download*” as seen when the Navy updated its Aegis combat software at sea in 2024 (Eckstein, 2024). This type of agility, performed as a normal course of events, means battlefield advantages can be gained (or lost) through fast code deployment, and gain rapid mission impact through multi-platform capabilities, placing a premium on the ability to test updates quickly and confidently.
- **AI-driven decision support**—AI helps process the flood of sensor and intelligence data to accelerate targeting, logistics, and command decisions. Machine-learning algorithms serve as copilots, offering recommendations or automating analysis for human commanders. In practice, AI-driven battle management might flag threats or suggest courses of action within seconds. The challenge for testers is validating that these AI systems are reliable and bias-free under the pressures of real operations. Ensuring that AI makes sound recommendations on a jammed, contested, chaotic environment—be it in grey-zone conflict or full-scale war—requires extensive scenario testing (Goode et al., 2025) and new T&E methods (e.g., red-teaming the AI's decision logic; Freeman et al., 2025).
- **Digital twins and simulation**—Militaries increasingly rely on high-fidelity software models—digital twins—to evaluate new concepts of operation, predict outcomes, and rehearse



missions through live-virtual-constructive environments. When continuously updated, these simulations can provide commanders with predictive insight by exploring “what-if” scenarios without real-world risk. The Pentagon envisions using such joint battlespace simulations to assess tactics and system integration in silico. However, their value depends entirely on fidelity. Test agencies must verify that digital twins accurately reflect operational reality, or they risk producing overly optimistic results that mislead decision-makers (AIAA ,202). Simulation is a powerful complement to live testing, but it requires rigorous validation and, given current technological and budget constraints, cannot fully replace real-world trials.

- **Reprogrammable weapons systems**—Missiles, drones, ships, and tanks are becoming software-reliant systems whose capabilities can shift through code rather than hardware changes. For example, a single UAS can move from reconnaissance to electronic attack by loading different software. Replicator swarms illustrate this flexibility: their behaviors and targeting priorities evolve through updates, not redesigns. This agility offers significant operational advantages but demands that each new release or configurable module truly adds value in the field. AI-augmented development accelerates reconfiguration and customization, increasing the need for tailored testing to ensure reliability. Testers must validate these rapidly changing systems under combat conditions, confirming that enhancements do not introduce instability or unintended interactions.
- **Autonomous and semi-autonomous systems**—UAS swarms, uncrewed surface/undersea and robotic ground vehicles, provide an expanding range of systems now operating with minimal human intervention, guided by software-based rules of engagement. These autonomous and semi-autonomous weapons adapt in real time to sensor inputs and shifting conditions, serving as powerful force multipliers but dramatically increasing testing demands. For example, how can testers ensure a drone distinguishes targets correctly in the fog of war, or that human–AI teams perform reliably under fire? Answering these questions requires new T&E methods and scenarios—live swarm exercises, virtual war-gaming, and stress tests across contested environments, such as GPS jamming, electronic warfare, and adversary deception. No test fully anticipates a wily opponent, so system architectures must also flex rapidly to incorporate responses when new behaviors emerge in the field.
- **Cyber operations as warfare**—Cyberspace is now a battlefield in its own right, where attacks on networks, satellites, and critical infrastructure can be as decisive as kinetic strikes. In software-defined warfare, adversaries may seek to defeat systems through hacking, spoofing, or software supply-chain compromise rather than physical destruction, driving the need for rigorous cybersecurity and resilience testing. Networked weapons and command systems must be hardened and evaluated against realistic cyber-attack scenarios. Modern operational testing increasingly incorporates red-team cyber assaults on systems, networks, and development pipelines alongside physical stress tests. To address these threats, the Pentagon is investing in advanced cyber ranges and simulations (Podnar et al., 2021) to assess system performance under sustained cyber pressure, recognizing that even the most advanced digital capabilities are ineffective if their software can be silently subverted.

Speed, accuracy and scale of new solutions to the fleet and force are necessary for victory (McNamara et al., 2025). Pursuing a strategy of software-defined warfare offers unprecedented adaptability and capability by leveraging software’s malleability. With this promise comes heightened complexity since success depends on whether these software-driven capabilities work as intended in the chaos of war. Each of the above tenets introduces new dimensions for T&E, from assuring AI decision aids, to validating sprawling simulations, to stress-testing autonomous behaviors and cyber defenses. The following sections examine how the Pentagon’s development and acquisition practices are evolving in this software-centric era,





and why rigorous, innovative, and independent testing is critical to realizing the potential of software-defined warfare.

## A Revolution in How the Pentagon Builds Technology

The U.S. military must advance beyond upgrading individual program-of-record systems—it now must reinvent how it conceives, develops, coordinates, and fields its arsenal. At the core of this concept is software-defined warfare: a reoriented focus on adaptability through code, powered through AI-driven decision-making, supported by modular upgrades, and continuous improvement in performance of forces. What once required decades of hardware cycles can now happen in months or even weeks through software-driven iteration. This section highlights two of the most dramatic shifts—AI-augmented development and accelerated acquisition—showing how they operationalize the tenets of software-defined warfare while creating urgent new challenges for T&E.

### AI-Augmented Development: Speeding the Software Cycle

A key driver of software-defined warfare is the integration of AI directly into the development process—an embodiment of software-defined warfare’s principle of *rapid adaptability*. Generative AI and machine-learning tools now help design systems, modernize legacy code, and even draft acquisition documents. The U.S. Army’s experimental “AI Flow” program, for instance, can generate an acquisition document 80% complete in minutes—a task that once took weeks (Fox, 2024). In 2023, the Pentagon underscored this shift with an \$800 million contract for advanced large language models (LLMs; CDAO, 2025). From automating code and documentation to shaping design decisions, AI copilots are accelerating delivery and embedding speed and flexibility into every stage of the cycle.

AI’s role is expanding beyond writing code to actively supporting testing and evaluation, as shown in Figure 2. Here, AI-driven tools extend the digital twin concept by using validated data to create simulations of complex systems and highlight vulnerabilities in real time. Early applications include automatically generating test cases from requirements, tracing system dependencies, and surfacing compliance gaps that might otherwise go unnoticed (Bain, 2025). By producing dozens of scenarios or code fixes rapidly, AI has the potential to accelerate delivery while catching subtle flaws that human testers could miss. This capability points toward a future where T&E itself becomes iterative and adaptive, evolving at the same speed as the software-defined warfare systems it safeguards.



**Figure 2. Applying Digital Twins to Accelerate Software Development and Testing**

Human experts (including testers) must therefore learn new workflows to collaborate with AI assistants, double-check AI outputs, and guard against failure modes introduced by AI-created

code or data. For instance, if an LLM recommends a code change, who ensures that the change is secure and reliable? The burden falls on T&E professionals to develop new verification frameworks suited to AI-generated software—frameworks that may involve algorithmic transparency, adversarial stress tests, and continuous monitoring (Schmidt, 2025). In short, AI has turbocharged the Pentagon's software cycle, aligning with software-defined warfare's tenet of rapid adaptability, but it also amplifies the imperative for agile testing at machine speed.

### **Accelerated Acquisition and Agile Deployment**

Another pillar of the Pentagon's software revolution is reforming how systems are acquired and deployed, treating individual platforms as components of larger, *reprogrammable capabilities* that span software-defined weapons and cyber warfare. Traditional acquisition cycles take years, yet modern capabilities—assembled from distributed subsystems across the battlespace—must evolve in hours, days, or weeks, creating a fundamental mismatch. To address this gap, the Pentagon has begun using flexible contracting authorities and agile methods to dramatically shorten acquisition timelines (AFLCMC, 2025).

For instance, program offices are using the authorities for using Other Transaction Authority agreements and Commercial Solutions Openings to bypass the usual red tape and reach non-traditional vendors (Gallium Solutions, 2025). Under Project Replicator, the Pentagon awarded contracts for new drone prototypes in as little as 110 days from concept to award, compared to the 12–18 months typical of older processes (Tucker, 2025). This pace matches the logic of software-defined warfare: field more adaptable, updatable systems that can be kept up to date faster than adversaries can counter them. It can be done today with significant effort for high priority mission capabilities—but the whole environment must be reoriented to delivering rapid capability at scale instead of systems.

Defense acquisition teams are also experimenting with LLM-based assistants and automation throughout the procurement and integration pipeline (Magnuson, 2025). By generating interface code, drafting requirements, and even allocating software tasks, these tools echo the *AI-driven decision support* described in the Software-Defined Warfare: Key Tenets section. Combined with iterative DevSecOps methods, they allow the Pentagon to push out software updates and modular hardware prototypes at a pace unthinkable a decade ago. This adaptability makes the force smarter and faster, but it also places immense pressure on T&E processes that must validate these updates in near real time.

Figure 3 depicts this transformation, where the mission systems hosted in combat platforms, from tanks to ships to UAS swarms, now receive over-the-air updates that alter their capabilities rapidly (Sherbinin & Gray, 2025). These AI-enabled systems not only generate tactics and recommend strategies but also deliver software patches that can reconfigure weapon systems at the speed of need, giving U.S. forces agility through code rather than costly and time-consuming hardware redesign. Figure 3 depicts software-defined warfare in action—an adaptable, reprogrammable force that can outpace adversaries by shifting faster than they can respond. Yet the very speed and scale that make this revolution possible also heighten the stakes: without rigorous T&E, rapid reprogramming risks introducing unseen vulnerabilities into the fight.





**Figure 3. An Agile Force is a Reprogrammable Force**

The surge in speed and quantity of new software-defined warfare systems underscores the Pentagon’s testing dilemma. The central question looms: how do we know these rapidly fielded, software-defined tools will actually work in the chaos of war? Contracting and delivering a system quickly is one thing; making sure it will work, with real operators, fully supported, in a tactical environment is another. As the next sections explore, the imbalance between rapid fielding and thorough testing is growing—and with it, the danger that America’s software-defined arsenal could prove brittle when it matters most.

## **Fielding Faster Than We Can Test**

This section examines the growing tension between the Pentagon’s accelerating software cycle and the slower pace of assurance. Software-defined warfare promises rapid adaptability, autonomy at scale, and reprogrammable weapons updated overnight—but each breakthrough multiplies the necessary action of verification and validation. Fielding cycles now measured in weeks outstrip test capacity, forcing greater reliance on models and new methods even as interdependent, software-updatable systems flood the force. The hardest challenge lies in proving interoperability and resilience across the joint “kill web”—autonomy, swarms, electronic warfare, cyber, and over-the-air updates—where failures often surface at the seams.

## **The Growing Imbalance Between Fielding and Testing**

Even as military arsenals evolve at accelerated speed, the T&E mechanisms needed to validate those new tools are struggling to keep up. According to a recent Defense Science Board (DSB) study, the gap between fielding and testing is not closing but widening, with potentially serious consequences (Evans et al., 2024). Officials and experts are increasingly worried about a mismatch where the Pentagon can deploy new software-defined systems in months but cannot test them thoroughly in time. The very tenet of *rapid adaptability* means capabilities can be reprogrammed at will, yet each change requires verification and an appreciation that other things didn’t break when a fix is introduced, which is increasingly hard to assure at the same pace. Accidental introduction of new bugs with software updates is a known issue and requires continuing focus on resilient architectures, automated processes, and regression testing to reduce the risks (DOT&E, 2024).

The DSB study also found that many emerging technologies “often require development of new analytical tools and testing capabilities for verifying and validating system performance,” and in some cases physical testing might not even be possible, forcing greater reliance on modeling and simulation study (Evans et al., 2024)] This reliance connects directly to the *digital*

*twin* concept of software-defined warfare. Models and simulators may allow updates to be stress-tested virtually, but unless they are themselves validated, they risk creating false confidence. This DSB study emphasized that to test leap-ahead technologies, such as hypersonic missiles, directed-energy weapons, or AI/ML systems, the Pentagon must invest in novel verification methods—and validate those methods themselves—to fill gaps where traditional live testing falls short, as discussed in the case study of the Joint Simulation Environment (JSF) shown on the previous page.

### ***Case Study of the Joint Simulation Environment.***

The F-35 Joint Simulation Environment (JSE) stands as both an essential achievement and a cautionary example. It ultimately delivered sophisticated, physics-based testing capability needed to evaluate combat scenarios that cannot be affordably or safely executed on an open range, as shown below.



However, its development demonstrated how hard it is to build a simulation environment that decision-makers can trust—one that accurately reflects real-world data, supports dynamic interaction, and produces results credible enough to inform combat-readiness decisions. Achieving this level of fidelity requires rigorous upfront planning and a clear understanding of the underlying physics, data dependencies, and system integration challenges.

The initial JSE effort underestimated this complexity and began with an incomplete vision of what a credible environment required (DOT&E, 2023), leading to abandonment and restart of the original design. The reconstituted JSE ultimately succeeded, enabling the F-35 program to reach full-rate production and establishing a gold standard for validated warfighting simulation (Naegele, 2024). However, it became operational only after more than 1,000 aircraft had already been produced under incremental acquisition authorities. Consequently, issues discovered during IOT&E—though correctable—were more expensive to resolve and introduced greater early operational risk than if identified earlier in the development life cycle (Guertin, 2022).

Without significant changes in technology, workforce, and oversight, imbalance between fielding and testing could undermine both safety and effectiveness. The reality today is that dozens of AI-driven drones, autonomous vehicles, and software applications are pouring into the joint force each year—an embodiment of the autonomy and reprogrammable weapons systems tenets. Each may be deployable, but can the Pentagon ensure that they are safe, effective, and interoperable? As another DSB study on applying AI in DoD missions starkly put it: if that trend continues, how can the Pentagon ensure that each system, much less the whole force, will work reliably under combat stress (DSB, 2023).

Unlike earlier eras when a single platform could be tested in isolation, today's systems-of-systems are woven into sprawling, interdependent networks. For example, a drone's effectiveness may rely on satellite feeds, AI-enabled command centers, and/or coordination with swarms of other autonomous assets. These dependencies are the essence of the Pentagon's Joint All-Domain Command and Control (JADC2) concept, which involves an intricate "kill web" linking sensors, shooters, and decision nodes across domains. As shown in Figure 4, unvalidated or untested interactions within this web can introduce hidden vulnerabilities or faulty



assumptions of behavior, creating risks that ripple through the entire network. Testing one component in isolation is no longer sufficient—testers will now need to understand the behavior of the entire web of connections in order to assess effectiveness, a challenge that grows exponentially with every new link.




**Figure 4. Unvalidated Connections Create Unseen Vulnerabilities in JADC2 Kill Webs**

The bottom line is that the U.S. military’s ability to rapidly field new software-defined capabilities now outpaces its capacity to test them under realistic conditions. This imbalance creates a risk of deploying systems that are not fully vetted for the rigors of combat. Recognizing this, defense analysts are emphasizing the need to cover the gap between the rate of fielding new systems and the capacity to test them thoroughly (Schmidt & Guertin, 2025).

### Testing Challenges in the Era of Kill Webs and Autonomy

Ensuring that a single new weapon works as designed is hard enough; ensuring that hundreds of new, software-defined systems all work together, each system also possibly being updated, in a dynamic battlespace is a monumental challenge that doesn’t scale with traditional testing methods. The Pentagon’s push for integrated, all-domain operations means that victory rests on cohesion and interoperability, e.g., Air Force systems must seamlessly coordinate with Army, Naval, Space Force, and allied systems in real time. This ambition reflects the *autonomy* and *cyber-as-warfare* tenets of software-defined warfare but also magnifies the difficulty of verification.

Recent experience underscores the difficulties. In one experiment involving Project Overmatch, a capability was fielded rapidly via over-the-air updates (a textbook example of *reprogrammable weapons systems*) but validating its interoperability at scale proved complex and time-consuming (DOT&E, 2024). Large-scale joint exercises and integrated tests—where multiple services stress their systems together—are logistically hard, expensive, and often cut for schedule or budget reasons. Under tight timelines, many rapid acquisition programs simply lack operationally realistic test environments, undermining the *rapid adaptability* that over-the-air updates were meant to provide.

Program offices or contractors, eager to deliver quickly, sometimes choose lower programmatic cost or schedule risk at the expense of higher risk-to-operational failure. This tendency is especially true for integration efforts and how systems interact with each other as a harbinger of broken schedules and unfulfilled promises. These shortcuts are particularly risky in

the context of software-defined warfare, where each system's function may depend on another's data or algorithms. The result is that systems reach deployable status with only limited understanding of their true performance in a system-of-systems context. Warfighters may be handed new drones, sensors, or AI tools that individually passed lab tests, but no one knows for sure if they will mesh when plugged into the joint kill web.

This challenge is not merely hypothetical—the gap between lab success and battlefield reality has already been exposed. Troops have fielded new technologies only to discover too late that a drone could not link with a Navy ship or that an AI guidance system collapsed under electronic jamming (Albon, 2025). As shown in Figure 5, swarms that appear flawless in controlled test environments can unravel amid the chaos of a contested battlefield, where cyberattacks and spectrum interference often degrade communications and coordination. Failures like incompatible radios and software crashes underscore the danger of testing *autonomy* and *AI-driven decision systems* only under ideal conditions. Building resilience for dozens of interconnected agents demands joint, all-domain trials at scale—exercises that remain rare despite their necessity.



**Figure 5. Swarm Success in the Lab, Chaos in the Fight: Why Testing Must Match Reality**

To their credit, Pentagon stakeholders are not blind to these challenges. There have been efforts to modernize testing itself—to “test the way we fight,” meaning more joint, all-domain evaluations in environments that mirror the drone-saturated, data-driven battlefields of tomorrow (DoD, 2024). The updated operational T&E guidance in DoDI 5000.98 aligns with the *digital twin and simulation* tenet, which envisions more advanced models and automated test tools for AI-enabled capabilities. National defense research teams are exploring how to close the “software understanding gap” (CISA, 2025) and the Armed Services are also investing in exercises like Project Convergence to identify interoperability issues early (South, 2022).

However, these efforts face steep institutional and resource hurdles. Testing in an all-domain, software-defined context demands not only new technology—simulators, distributed ranges, AI-driven test orchestration—but also sufficient skilled people and resources. Defense analysts are observing a troubling inconsistency: at the very moment such broad, innovative testing is most needed (Schmidt & Guertin, 2025). In addition, the Pentagon has dramatically downsized the very office responsible for independent operational test and evaluation across the services (Hegseth, 2025). This irony underscores the theme of this paper: the Pentagon is

charging into the era of software-defined warfare, but unless it invests equally in testing and oversight, speed could become fragility.

## **Trust, Technology, and the Testing Tightrope**

This section explores the paradox at the heart of modern T&E: the Pentagon is best served by preserving independent oversight while also embracing new simulation and automation tools. The stakes high in this era of software-defined warfare, where systems are rapidly adaptable, digitally twinned, reprogrammable, autonomous, and cyber-contested. Unlike traditional weapons platforms, software-reliant systems can change rapidly through updates that introduce new features *and* new risks, making independent T&E critical to operational safety and trust. Without adequate scrutiny, there is a real danger of deploying unproven code or brittle autonomy into contested battlespaces, where the margin for error is razor-thin. This section also examines the promise and peril of digital twins, AI-driven test agents, and large-scale simulation, emphasizing how these tools can expand coverage but are not yet replacements for seasoned testers.

## **The Urgency to Modernize Independent Test and Evaluation**

For decades, verifying the adequacy of test plans and the validity of their results has consistently demonstrated its value. The Director of Operational Test & Evaluation (DOT&E)'s annual reports are replete with examples of programs falling short of delivering needed capability. An independent evaluator of T&E performance and data analysis has proven value in supporting software-defined warfare. An organization like DOT&E and the affiliated Service's operational test agencies ensure that teams responsible for product development are held to high standards for how tests are written and how the results are analyzed.

New approaches are also needed to subject teams developing new weapons systems—and major upgrades to legacy systems—to rigorous external validation to ensure their products are tested under realistic combat conditions and evaluated objectively. This role is essential in the context of software-defined warfare because code can be changed overnight, and the challenges of not having a test plan and the resources to execute it risk deploying unverified updates into combat. In essence, DOT&E's job is to ensure weapons systems are “battle-tested before battle,” identifying critical flaws that must be fixed before warfighters' lives are at stake.

Realistic testing is indispensable in an era of *AI-driven decision support* and *autonomous systems*, where critical failure modes often emerge only in contested, unpredictable environments. Meeting this challenge requires enterprise-level resources to design and orchestrate rigorous field events and multi-service exercises. These integrated, cross-domain testing are needed to ensure true interoperability and realize the kill-web vision of software-defined conflict.

Modern testing must validate not just a single platform in isolation, but the complex web of interactions across satellites, ships, aircraft, drones, and ground systems that now define how wars are fought. As shown in Figure 6, illuminating and closing gaps in kill webs is needed to highlight untested links, validate cross-domain connections, and ensure interoperability under real-world stresses. Without this type of holistic system validation, many critical vulnerabilities may remain hidden, e.g., a contractor may certify a drone as mission-ready yet fail to test it under electronic jamming or in coordination with allied assets—precisely the kinds of weaknesses adversaries exploit in cyber-as-warfare environments. Oversight ensures that speed and innovation do not outpace resilience and trustworthiness.

Given the increasing complexity of software-defined, networked systems, independent oversight of testing is more crucial than ever. Yet recently resources for advancing the practice of T&E have been downsized significantly (Hegseth, 2025). However, reducing oversight before





those methods and tools mature runs counter to the very principles of *responsible autonomy* and *digital twin validation* (Fredenburg, 2025).



**Figure 6. DOT&E Helps Validate the Webs of War**

In the context of software-defined warfare, independent comprehensive testing provides something intangible yet invaluable: *trust*. When DOT&E signs off on a system after rigorous evaluation, military leaders and troops can have confidence that T&E professionals tried hard to break the system before it reached them. That trust is hard to quantify, but easy to appreciate when lives are on the line. If independent oversight is weakened, the Pentagon risks succumbing to optimistic bias or deploying unvetted updates. In a software-defined era, DOT&E is not just an overseer—it is the steward of confidence.

### **Embracing Simulation and Automation: Opportunities and Risks**

Facing an explosion of complex, software-defined systems to test (and with fewer human testers on hand), the Pentagon is understandably looking to technology for help. Advanced simulation environments, digital twins, and automated test tools embody the tenets of digital twins and simulation and rapid adaptability. If implemented properly, they could help bridge the gap between rapid fielding and rigorous evaluation, allowing the Pentagon to scale its testing to match the speed of its software-defined arsenal (Software Engineering Institute [SEI], 2020).

For example, instead of conducting dozens of live-fire exercises to ensure a new UAS works in multiple scenarios, testers could run thousands of virtual trials in a high-fidelity simulator. Such tools could evaluate interoperability by linking virtual instances of an Air Force command system and an Army missile battery. Automated cyber agents could probe vulnerabilities, reflecting the cyber warfare domain, while digital twins could predict system performance after software updates. In theory, these methods could vastly expand what testers can cover, especially with limited time and resources.

However, current reality has not caught up to this vision. The advanced test automation needed for realistic testing remains at relatively low technology readiness levels—essentially prototypes themselves (Guertin, 2022). Many tools remain too immature for large-scale use, a serious liability for reprogrammable weapons that require continuous validation. Without robust research and development (R&D) investment, the promise of trustworthy *simulation* and *automation* remains out of reach.



In practice, building a credible digital twin of a complex weapon system is hard. High-fidelity models demand massive datasets, specialized expertise, and significant computing power to achieve accuracy (Hahn et al., 2023). Even small modeling errors can yield misleading “rosy outcomes” that collapse under combat stress—a risk that is especially acute for *AI-driven decision support*, where flawed assumptions in simulation can translate into catastrophic real-world failures. As a result, the Pentagon cannot simply simulate its way out of live testing until confidence in these models improves substantially.

The over-reliance on automation without human judgment is another risk. While a simulation may appear to provide near-perfect stability and success, the actual reality under combat stress may be quite different, e.g., systems failing, swarms unraveling, and missions collapsing, as shown in Figure 7. In practice, testing is not a simple box-checking exercise. Instead, intuition, creativity, and skepticism are needed that only experienced testers bring. If seasoned professionals are replaced too quickly with immature AI tools, the Pentagon risks gaining neither trusted human insight nor validated digital capability. This danger is especially perilous with *autonomous systems*, whose unpredictable behaviors cannot be fully captured by models. Test tools themselves must therefore be tested and accredited as rigorously as the systems they claim to evaluate.




**Figure 7. Simulation vs. Reality—the Cost of Over-Reliance**

The Pentagon leadership recognizes both the promise and pitfalls here. They argue that a leaner workforce forces innovation—indeed Pentagon leadership and the Armed Services are experimenting with analytics and automated testing to cover more ground with fewer people. Going forward, the most likely outcome is a hybrid approach: human-led design and analysis paired with machine-speed execution of simulations. This hybrid aligns with the human-machine teaming ethos already central to software-defined warfare.

For instance, an AI-driven orchestration system could coordinate distributed simulations across labs and crunch terabytes of test data, spotting anomalies faster than humans alone. This approach could allow small test teams to supervise continuous, large-scale war-gaming of our systems versus simulated adversaries. However, achieving this vision requires both patience and investment. The technologies needed to automate operational testing at scale are still nascent and cutting human testers before it matures risks creating oversight gaps. Thus, replacing seasoned testers faster than reliable automated tools can be fielded is not a recipe for confidence.

In summary, advanced simulations and test automation will be indispensable for tackling the complexity of software-defined warfare. They embody its tenets by offering scalability, adaptability, and speed. Yet in 2025 they remain supplements to—not substitutes for—experienced testers and realistic trials. The Pentagon must walk a fine line: pursuing new

technologies aggressively while acknowledging their current limits. To succeed, Pentagon leaders must remember a core principle: in software-defined warfare, agility without trust is fragility. The next section explains why getting this balance right is so critical, by examining the risks of under-tested systems in the field.

## Risks of Under-Tested Warfare

Modern U.S. military strategy assumes that our high-tech weapons will work as intended when called upon. If that assumption fails due to inadequate testing, the consequences could be dire. This section examines several distinct risks that can arise if the military embraces software-defined warfare without sufficiently comprehensive T&E.

### Integration Failures in the Kill Web

One risk is that the fragility of integration—the chance that our mosaic of integrated systems needed to support the joint force will not mesh into the coherent kill web envisioned by strategists and war planners. Software-defined warfare thrives on *AI-driven decision support* and *autonomous systems* coordinating across domains. Without robust joint T&E, however, there is no assurance that this myriad of disparate platforms will interact seamlessly.

Figure 8 depicts a central danger in software-defined warfare: peer adversaries won't strike at our strongest systems, but at the fragile seams between them. The Navy ship on one side and the Army missile battery on the other symbolize powerful platforms, yet a hacker gaining entry to the network underscores how integration fragility can undermine them both. AI-driven decision support and cross-domain coordination promise a coherent kill web. If systems are validated only in isolation, however, hidden incompatibilities persist. Without robust joint testing these inter-connections will remain brittle and unseen, exposing vulnerabilities that adversaries are poised to exploit, creating vulnerabilities at the seams.




**Figure 8. Adversaries Don't Attack the Strongest Links, They Exploit the Weakest Seams**

For example, an adversary exploiting a single untested interface between two Services' systems could turn a seemingly hardened network into a battlefield liability. In an era of software-defined warfare—where capabilities are stitched together dynamically and updated at machine speed—only rigorous, end-to-end joint testing can reveal hidden interdependencies that determine whether kill webs hold or unravel under pressure. Cross-service interoperability trials and large-scale exercises are thus essential, despite their cost and complexity, are indispensable. Without them, the elegant mosaic of interconnected systems we envision and rely

upon risks becoming a brittle patchwork—one that could fracture in the fog and fury of actual conflict.

### False Confidence in Untested Systems

Another risk is false confidence—believing a new capability is a significant improvement but later uncover Achilles’ heels that only rigorous testing would have exposed. The ethos of rapid adaptability can seduce decision-makers into fielding systems based on demos or contractor reports that fail to reveal the whole story. History is full of examples of weapons that looked promising in controlled trials but stumbled in combat (Wilkinson, 2014).

Figure 9 highlights the illusion of being “ready” in the software-defined warfare era. A swarm of drones may appear mission-capable in controlled tests, but the real proving ground comes when adversaries introduce stressors like jamming, GPS spoofing, or hidden software bugs—the very threats depicted beneath the shield. What appears resilient in peacetime can quickly unravel under hostile conditions, exposing cracks in autonomy and code. Without rigorous testing across these edge cases, the confidence inspired by a glowing “ready” status is dangerously misleading. These types of failures quickly expose the fragility of *autonomous and reprogrammable systems* when not fully tested across edge cases.




**Figure 9. Stress Reveals Hidden Flaws**

DOT&E reports are full of such “unknown unknowns,” uncovered only during rigorous, realistic testing (DOT&E, 2024). Runtime monitoring and safety guards can reduce these risks, but without a safety net, flaws remain hidden until the battlefield, where they can cause catastrophic consequences. Skipping or truncating operational testing breeds the illusion of readiness—leading commanders and troops to assume systems are reliable when they are not. In software-defined warfare, such misplaced confidence is more dangerous than no capability at all, because it encourages faulty tactical decisions built on brittle foundations.

### Eroded Trust and Adoption Delays

Paradoxically, cutting corners on testing can slow down adoption of new technology in the field. Software-defined warfare depends on *human-machine teaming* and trust in *AI-driven decision support*. Yet frontline commanders and troops may hesitate to trust unproven systems with their lives. Trust is built via transparency and rigorous vetting: soldiers gain confidence when they know a system has been tested thoroughly and when they themselves understand its strengths and failure modes.

If a new drone or AI assistant arrives with scant evidence of its performance in contested conditions, operators will naturally be cautious. In well-tested programs, leaders receive detailed reports on what worked, what failed, and what to watch for—establishing a baseline of trust. Without that vetting, doubts linger. Troops might receive powerful new tools but use them



gingerly or not at all, fearing they could backfire, especially for *autonomous decision-making* or *cyber-sensitive systems*.

Figure 10 shows the difference between rushing untested systems into the field versus rigorously validating them before deployment. In the top panel, rushed fielding leaves troops untrained, creates uncertainty, and leads to operational pauses as confidence erodes. Warfighters hesitate to use equipment they cannot trust, undermining the very purpose of *rapid adaptability*. In contrast, the bottom panel shows how rigorous testing earns a DOT&E “approved” stamp, building trust and ensuring systems are combat ready. With confidence in performance instilled by T&E under realistic conditions, warfighter adoption accelerates, and deployment becomes more effective. Speed ultimately flows from trust—and trust in software-defined warfare depends on testing that provides evidence, not assumptions.



**Figure 10. Rushed Fielding Yields Delays; Rigorous Testing Builds Trust and Accelerates Deployment**

### **Tactical and Strategic Surprise**

It is also essential to consider how the lack of rigorous testing provides adversaries with opportunities to exploit and degrade our capabilities. Our near-peer competitors relentlessly probe for weak links in our digital arsenal. A flaw in a cyber-dependent or software-defined system that slips through testing is a flaw our enemies will eagerly exploit. Figure 11 shows how a single exploited weakness like electronic warfare jamming can ripple throughout an entire kill web. Adversaries constantly search for these fault lines in our networked joint force. A vulnerability in a satellite relay, a radar node, or software in a reprogrammable drone can cascade throughout the system, degrading every connected platform. The lack of rigorous testing doesn’t just risk isolated failures—it hands our competitors the opportunity to unravel our whole web of advantage.



**Figure 11. A Single Point of Failure Can Unravel an Entire Kill Web**

Thorough testing is how we harden these systems, so our adversaries do not get easy openings. By neglecting testing, we risk strategic surprise. The United States could enter a conflict expecting its software-defined arsenal to dominate, only to see it neutralized—or even turned against us—due to an unvetted defect. The very malleability that makes software-defined warfare powerful also makes it fragile if oversight lags.

In sum, a culture of “field it fast and fix it later” increases operational and safety risk significantly. In contrast, a culture that values rigorous OT&E forces the Pentagon to confront inconvenient truths early, denying adversaries an easy win. The need for rigorous OT&E is especially true in *autonomous and AI-enabled systems*, which adversaries are eager to jam, spoof, or deceive. Only testing across adversary-relevant scenarios can prevent our strengths from becoming liabilities.

Collectively, these risks paint a stark picture: under-tested software yields high-risk warfare. Speed without assurance gambles with lives and strategy alike. The good news is that the Pentagon need not choose between agility and rigor. With the right reforms and investments, it can have both—leveraging the tenets of software-defined warfare for rapid innovation while ensuring those capabilities are battle-ready through robust validation. The next section offers recommendations to achieve this balance and build a resilient test enterprise fit for the digital battlespace.

## **Recommendations: Building a Test Enterprise for Software-Defined Warfare**

To unlock the full potential of software-defined warfare—and avoid the vulnerabilities described in the previous section—the Pentagon must reimagine its T&E enterprise. As shown in Figure 12, this reimagining requires an integrated workflow that combines simulation and digital twins, continuous testing, AI-enabled automation, a collaborative workforce, and a robust and properly resourced DOT&E office that is organizationally independent of systems under test. Together, these elements will form a resilient test ecosystem capable of validating systems at the speed and complexity demanded by modern conflict. This section presents key recommendations for building an enterprise T&E infrastructure and workforce that can support these capabilities effectively in the context of software-defined warfare.





**Figure 12. Key Elements in the T&E Enterprise for Software-Defined Warfare**

### **Invest in Advanced Simulation and Digital Twin Infrastructure**

The Pentagon should expand and connect its test ranges and labs into a persistent digital proving ground. This investment is central to the *digital twins and simulation* tenet of software-defined warfare, which envisions predictive models that evolve in parallel with production weapons systems. Funding should prioritize high-fidelity simulators and digital twin models for complex systems, networked together to replicate entire kill webs.

For example, a joint simulation might concurrently model an Air Force drone swarm, an Army missile battery, and a Navy command ship to evaluate end-to-end performance. Such environments could run thousands of virtual scenarios, complementing physical tests and stress-testing AI-driven decision support at scale. To ensure credibility, however, validation must be prioritized: digital twins should be compared rigorously to live operational data. A concerted R&D effort—through DARPA, the Test Resource Management Center, and Service labs—is needed to mature these tools from experimental prototypes into trusted T&E assets.

### **Integrate Continuous Testing into DevSecOps Pipelines**

Taking a cue from modern software practice, the Pentagon should adopt a mindset of testing early, often, and automatically across a system’s life cycle. An example is using DevSecOps pipelines, an integral part of the preferred Software Acquisition Pathway (Harper, 2025) to automate testing in incremental updates (SEI, 2025) to enable continuous testing as opposed to the testing at the end mindset. This approach directly supports the rapid adaptability and reprogrammable weapons systems tenets, where updates can be deployed quickly and verified immediately. Continuous testing requires embedding automated test suites in DevSecOps pipelines for both software and hardware-in-the-loop. Every new code commit or modular upgrade should trigger regression checks and cybersecurity scans. Whenever feasible, prototypes and updates should be fielded in controlled “beta” capacity that can be rolled back quickly if deviant behavior emerges, capturing operational feedback in near-real time.

An illustrative example would be continuously integrating and testing incremental updates to an autonomous UAS in a sandbox with warfighter participation, rather than deferring all validation to a single T&E event. By the time a capability is fully fielded, it will have already proven itself in multiple iterative trials—ensuring trust in rapid adaptability while reducing risk. This approach would allow operators and warfighters to stress-test the UAS under varied conditions, from electronic warfare interference to extreme weather, surfacing potential weaknesses long before combat deployment.

## **Leverage AI and Automation as Force Multipliers in Testing**

The T&E community should embrace AI tools not as replacements for humans but as amplifiers of their effectiveness. This strategy aligns with the *AI-driven decision support* and *autonomy* tenets of software-defined warfare, where AI agents work alongside humans to manage complexity at scale. AI can rapidly generate test cases, explore edge-condition combinations, and orchestrate distributed test events that would be infeasible to conduct manually (Schmidt, 2025). Generative AI could simulate adversary cyber tactics or electronic jamming against a system under test, exposing vulnerabilities in cyber-as-warfare contexts (Freeman et al., 2025). Machine learning algorithms can process massive test datasets, flag anomalies, and highlight performance trends invisible to humans.

Intelligent orchestration tools could manage networks of test assets—labs, ranges, and virtual environments—running large-scale scenarios overnight while humans review flagged issues in the morning. To realize this vision, the Pentagon should sponsor pilot projects in areas like automated test plan generation, adversarial scenario creation, and anomaly detection. This AI-augmented testing will enable the T&E workforce to scale coverage and focus on higher-order analysis.

## **Bolster T&E Expertise and a Collaborative Workforce**

Technology alone cannot deliver trusted outcomes; people are as essential as machines in the testing enterprise. The human-machine teaming dimension of software-defined warfare requires testers who understand both the software driving modern weapons systems and the operational contexts in which those systems fight. The shortage of experienced testers highlights the need to rebuild this expertise. The Pentagon should prioritize hiring and developing “cyber-warrior testers” with fluency in coding, AI, and mission operations. Recruiting talent from industry, coupled with new military and civilian training pipelines, will accelerate this shift. To capture the wisdom of experienced T&E professionals, the DoD should create cross-agency knowledge exchange programs.

For example, joint T&E task forces could bring together the Service’s test and evaluation agencies in the Army Test and Evaluation Command, the Navy’s Operational Test and Evaluation Force, the Air Force Test and Evaluation Command, and the Joint Interoperability Test Command, the non-Government trusted agents in our research and development teams (e.g. federally funded research and development centers and university applied research centers, as well as other deep experts in academia to tackle high-priority challenges collaboratively. Regular T&E summits and workshops can ensure new programs do not start from scratch when grappling with AI-powered systems. In short, building networks of T&E expertise will ensure the workforce tests not just hardware, but reprogrammable, autonomous, and cyber-contested systems, as well.

## **Maintain Rigor and Independence in T&E**

Finally, the Pentagon must preserve the principle of independent and comprehensive T&E. In the era of software-defined warfare where systems evolve continuously, independent oversight provides the trust anchor for every update. The office of DOT&E—or an equivalent—must retain authority to provide unvarnished assessments, free from political or programmatic pressures. Independent T&E should be reframed not as a bureaucratic brake but as an enabler of speed: by having the discipline of catching problems early, thus preventing catastrophic failures that derail campaigns.

Pentagon leadership must ensure that any workforce reductions are balanced by investments in tools and training. Likewise, Congress should require programs to allocate sufficient time and resources for OT&E events, even if that delays initial fielding. In practice, this process means applying a new rule of thumb for software-defined warfare: “Never trust an



algorithm in battle until it's been stressed in every way feasible." Continuous capability evolution is a strength only if each iteration earns confidence through testing. Independent OT&E ensures that every reprogrammable, AI-enabled system is not just innovative—but dependable.

By pursuing the above initiatives, the Pentagon can build a T&E infrastructure and culture that matches the flexibility and speed of software-defined warfare. The goal is a virtuous cycle where rapid innovation is met with equally agile verification, enabling the United States to confidently field cutting-edge capabilities that not only work as promised but can be trusted in the chaos of war.

## Concluding Remarks

Ironically, the Pentagon's digital revolution and its testing crisis are two sides of the same coin—both born from the dizzying pace of technological change. The U.S. military must innovate rapidly to keep its edge, embracing rapid adaptability through software updates, AI-driven decision support to manage complexity, autonomous systems to expand force projection, and cyber-as-warfare tools to contest new domains. Yet technology alone does not win wars; only technology that works when and as needed does. Ensuring that outcome depends on thorough and independent comprehensive T&E to verify that each tenet of software-defined warfare is reliable in the chaos of combat.

In their rush to field the future, defense leaders cannot afford to abandon the proven principles that undergird past successes: realistic training, iterative improvement, and honest evaluation of what works and what fails. There's an old military adage: "Never send a weapon into war until you've seen it fight." In the context of software-defined warfare, that adage must be updated to: "Never deploy an algorithm, a digital twin, or a reprogrammable system to battle until you've stressed it in every way feasible." Achieving that standard will require reinvention and sustained investment in the T&E enterprise—not as an optional safeguard, but as the foundation for realizing the promise of an AI-enabled, software-driven arsenal.

As the Pentagon races ahead with autonomous UAS swarms, AI-powered battle managers, and Internet-of-Things-style weaponry, it must continually ask itself: Who is testing these new technologies, and to what standard? Who ensures that our autonomous systems behave predictably under duress, that our AI-driven decision support resists bias (or at least, incur biases we understand), and that our reprogrammable weapons adapt without introducing new vulnerabilities? The coming years will be telling. We will see whether the United States can balance speed with safety, and innovation with validation, in the realm of defense.

In fact, this piece is all about propelling the kinds of change in outcome the new defense acquisition policy states as an objective. Ultimately, our men and women in uniform deserve to go into the next conflict with systems we know will have their backs—not with question marks. To get there, the Pentagon must rebuild a resilient testing safety net, updated for the digital battlespace. The future of American warfare will depend not just on how fast we can innovate, but also on how well we can test, validate, and trust what we create.

## References

- Air Force Life Cycle Management Center. (2025, March 10). *Modern software acquisition to speed delivery, boost warfighter lethality*.  
<https://www.afmc.af.mil/NEWS/Article/4115238/modern-software-acquisition-to-speed-delivery-boost-warfighter-lethality/>
- American Institute of Aeronautics and Astronautics (AIAA), & Aerospace Industries Association (AIA). (2020). *Digital twin: Definition & value—An AIAA and AIA position paper*.



- [https://www.aiaa.org/docs/default-source/uploadedfiles/advocacy/key-issues/digital-engineering/aiaa-ai-digital-twin-position-paper-\(nov-2020\).pdf](https://www.aiaa.org/docs/default-source/uploadedfiles/advocacy/key-issues/digital-engineering/aiaa-ai-digital-twin-position-paper-(nov-2020).pdf)
- Albon, C. (2025, July 15). *Jammed and confused: Alaska trial shows pitfalls of fielding US drones*. Defense News. <https://www.defensenews.com/pentagon/2025/07/15/jammed-and-confused-alaska-trial-shows-pitfalls-of-fielding-us-drones>
- Bain, W. (2025, June 11). *Real-time digital twins with AI/ML: A new level of battlefield intelligence*. Military Embedded Systems. <https://militaryembedded.com/ai/machine-learning/real-time-digital-twins-with-ai-ml-a-new-level-of-battlefield-intelligence>
- Chief Digital & Artificial Intelligence Office. (2025, July 14). *CDAO announces partnerships with frontier AI companies to address national security mission areas*. <https://www.ai.mil/Latest/News-Press/PR-View/Article/4242822/cdao-announces-partnerships-with-frontier-ai-companies-to-address-national-security-mission-areas>
- CISA. (2025, January 16). *Closing the software understanding gap*. Cybersecurity and Infrastructure Security Agency. <https://www.cisa.gov/sites/default/files/2025-01/joint-guidance-closing-the-software-understanding-gap-508c.pdf>
- Clark, B., Vaddi, D., & Sloman, H. (2023). *From sensor to shooter: Modernizing the kill chain for the era of software-defined warfare*. Hudson Institute. <https://www.hudson.org/national-security-defense/sensor-shooter-modernizing-kill-chain-era-software-defined-warfare>
- Defense Science Board. (2023, July). *Designing and deploying artificial intelligence for DoD missions*. U.S. Department of Defense. [https://dsb.cto.mil/reports/2020s/AI\\_Study\\_Final.pdf](https://dsb.cto.mil/reports/2020s/AI_Study_Final.pdf)
- Director, Operational Test & Evaluation. (2024, January). *FY2023 annual report*. U.S. Department of Defense. <https://www.dote.osd.mil/Annual-Report>
- DoD. (2024, December 9). *DoD instruction 5000.98: Operational test and evaluation and live fire test and evaluation*. <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500098p.PDF>
- DoD. (2025). *Software modernization implementation plan FY25–26*. <https://dodcio.defense.gov/Portals/0/Documents/Library/SW-Mod-I-Plan25-26.pdf>
- DoW. (2026, January 6). *Transforming the defense innovation ecosystem to accelerate warfighting advantage* [Memorandum for Senior Pentagon Leadership]. <https://media.defense.gov/2026/Jan/12/2003855657/-1/-1/0/TRANSFORMING-THE-DEFENSE-INNOVATION-ECOSYSTEM-TO-ACCELERATE-WARFIGHTING-ADVANTAGE.PDF>
- Eckstein, M. (2024, April 3). *US Navy making Aegis updates, training changes based on Houthi attacks*. Defense News. <https://www.defensenews.com/naval/2024/03/21/us-navy-making-aegis-updates-training-changes-based-on-houthi-attacks>
- Evans, E. D., Van Wie, D., Green, J. (eds.). (2024). *Report of the defense science board task force on test & evaluation*. Defense Science Board. [https://dsb.cto.mil/wp-content/uploads/2024/08/DSB\\_TE-Report\\_UNCLASS\\_FINAL\\_August-2024\\_Stamped.pdf](https://dsb.cto.mil/wp-content/uploads/2024/08/DSB_TE-Report_UNCLASS_FINAL_August-2024_Stamped.pdf)
- Fox, A. (2024, October 29). *A.I. Flow: Pioneering a new approach to artificial intelligence*. U.S. Army. [https://www.army.mil/article/280932/a\\_i\\_flow\\_pioneering\\_a\\_new\\_approach\\_to\\_artificial\\_intelligence](https://www.army.mil/article/280932/a_i_flow_pioneering_a_new_approach_to_artificial_intelligence)
- Fredenburg, M. (2025, August 4). *Cutting the DoD's testing office won't reverse US military decline*. The National Interest. <https://nationalinterest.org/feature/cutting-the-dods-testing-office-wont-reverse-us-military-decline>
- Freeman, L., Robert, J., & Wojton, H. (2025, June 24). *The impact of generative AI on test & evaluation: Challenges and opportunities*. *FSE Companion '25: Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering, Association for Computing Machinery (ACM)*. <https://doi.org/10.1145/3696630.3728723>





- Gallium Solutions. (2025, June 19). *What is an OTA contract? A beginner's guide for innovators and startups*. <https://www.galliumsolutions.co/post/what-is-an-ota-contract>
- Goode, W. S., Costello, R., & Shakarian, P. (2025). *Acquiring generative artificial intelligence to improve DoD influence activities* (RAND Report RR-A3157-1). RAND Corporation. [https://www.rand.org/pubs/research\\_reports/RR-A3157-1.html](https://www.rand.org/pubs/research_reports/RR-A3157-1.html)
- Guertin, N. H. (2022, July). *Digital twin assessment, agile verification processes, and virtualization technology report*. Office of the Director, Operational Test & Evaluation (DOT&E). [https://www.dote.osd.mil/Portals/97/pub/reports/\(U\)%20Digital%20Twin%20Agile%20Verification%20and%20Virtualization%20Technology%20Report.pdf](https://www.dote.osd.mil/Portals/97/pub/reports/(U)%20Digital%20Twin%20Agile%20Verification%20and%20Virtualization%20Technology%20Report.pdf)
- Hahn, A. S., deCastro, J., Tanaka, M., Lamb, C. (2023, October). *Evaluation of digital twin modeling and simulation* (SAND 2023-09884). Sandia National Laboratories. [https://www.sandia.gov/app/uploads/sites/273/2024/11/SAND\\_Digital\\_Twins\\_Final.pdf](https://www.sandia.gov/app/uploads/sites/273/2024/11/SAND_Digital_Twins_Final.pdf)
- Harper, J. (2025, March 7). *Hegseth issues edict on DOD software acquisition*. Defense Scoop. <https://defensescoop.com/2025/03/07/hegseth-memo-dod-software-acquisition-pathway-cso-ota>
- Hegseth, P. (2025, May 27). *Memorandum directing reorganization of the Office of the Director of Operational Test and Evaluation*. U.S. Department of Defense. <https://media.defense.gov/2025/May/28/2003725153/-1/-1/1/MEMORANDUM-DIRECTING-REORGANIZATION-OF-THE-OFFICE-OF-THE-DIRECTOR-OF-OPERATIONAL-TEST-AND-EVALUATION.PDF>
- Magnuson, S. (2025, February 27). *Generative AI used to speed up defense acquisitions*. *National Defense Magazine*. <https://www.nationaldefensemagazine.org/articles/2025/2/27/generative-ai-used-to-speed-up-defense-acquisitions>
- McNamara, W. M., Modigliani, P., & Nurkin, T. (2025). *Commission on software-defined warfare: Final report*. Atlantic Council, Washington, DC.
- Mulchandani, N. & Shanahan, J. N. T. (2022, September). *Software-defined warfare: Architecting the DOD's transition to the digital age*. <https://www.csis.org/analysis/software-defined-warfare-architecting-dods-transition-digital-age>
- Naegele, T. (2024, December 9). *JSE: How Air Force aims to get more pilots into world's best F-35 simulator*. *Air & Space Forces Magazine*. <https://www.airandspaceforces.com/jse-air-force-f-35-simulator-more-pilots/>
- Podnar, T., Dobson, G., Updyke, D., & Reed, W. (2021, May 19). *Foundation of cyber ranges* (SEI Report CMU/SEI-2021-TR-001). <https://doi.org/10.1184/R1/13557566>
- Schmidt, D. C. (2025). *Software testing in the generative artificial intelligence era: A practitioner's playbook*. *IEEE Computer*, 58(7), 147–152.
- Schmidt, D. C., & Guertin, N. H. (2025, September 3). *The Pentagon's software revolution and its testing dilemma*. *War on the Rocks*. <https://www.warontherocks.com/2025/09/the-pentagons-software-revolution-and-its-testing-dilemma>
- Software Engineering Institute. (2020). *TwinOps combines digital twins and DevOps for better cyber-physical systems*. <https://www.sei.cmu.edu/annual-reviews/2020-year-in-review/twinops-combines-digital-twins-and-devops-for-better-cyber-physical-systems>
- Software Engineering Institute. (2025, May 12). *SEI study details ongoing DevSecOps adoption in DoD*. Carnegie Mellon University. <https://www.sei.cmu.edu/news-events/news/article.cfm?assetId=109121>
- Sherbinin, A., & Gray, A. (2025, February 10). *Software-defined warships: The Navy's digital future of necessity*. *War on the Rocks*. <https://warontherocks.com/2025/02/software-defined-warships-the-navys-digital-future-of-necessity>





- South, T. (2022, November 4). *Army's project convergence links sensors, shooters in joint experiments*. Army Times. <https://www.armytimes.com/news/your-army/2022/11/04/armys-project-convergence-links-sensors-shooters-in-joint-experiments>
- Tucker, P. (2025, March 7). *Pentagon aims to accelerate acquisition of new tech through software-contracting change*. Defense One. <https://www.defenseone.com/technology/2025/03/pentagon-aims-accelerate-acquisition-new-tech-through-software-contracting-change/403598>
- Wilkinson, S. (2014, February 26). *10 of history's worst military weapons*. HistoryNet. <https://www.historynet.com/10-of-historys-worst-weapons>



# Digital Engineering to Improve the Quality of a Capability Needs Statement

**Alfred Schenker**—works in the Software Engineering Institute's (SEI's) Software Solutions Division, working there for over 25 years. He works to improve software acquisition and product development practices throughout the armed services and other organizations. He has actively worked in software process, architecture, model-based systems engineering, and metrics. Before joining the SEI, Mr. Schenker spent over 20 years in industry as an active contributor in all phases of product development. Mr. Schenker is also an inventor and has obtained patents for a pressure switch (used in automotive airbag applications) and a manufacturing process to seal gas inside a vessel. [ars@sei.cmu.edu]

## Abstract

This paper describes how an acquisition program on the Software Pathway (SWP) can elaborate its software Capability Needs Statement (CNS) within a modern digital engineering environment using Model-Based System Engineering (MBSE) principles. This elaboration (i.e., view) transforms the CNS from a static description into an authoritative, model-based view of mission deficiencies, required enhancements, interfaces, and interoperability. This view enables sponsors and users to prioritize and validate their needs iteratively with traceable evidence. This approach will lead to a more defensible and streamlined software development strategy. It also enables the program to review the capability with stakeholder(s) to ensure adequate validation and develop a quality artifact that can be used for negotiating with contractors.

**Keywords:** Software Acquisition Pathway, Model-Based Systems Engineering, Capability Needs Statement, Model-Based Acquisition, Agile, Minimum Viable Product, Digital Engineering

## Introduction

This paper introduces clear rationale for an acquisition program on the Software Pathway (SWP) to elaborate its software Capability Needs Statement (CNS) within a modern digital engineering environment. This is accomplished by employing Model-Based System Engineering (MBSE) principles that transform the CNS from a static description into an authoritative, model-based view of mission deficiencies, required enhancements, interfaces, and interoperability. This view enables sponsors and users to prioritize and validate their needs iteratively with traceable evidence rather than with documents. The motivation for developing this view is to achieve a digital representation of the system. This representation is available for the express purpose of capability definition, elaboration, and validation, and provides the acquirer with a superior mechanism for independent verification and validation. It also provides the user with a higher likelihood that the technical solutions for the capabilities will better meet their needs.

We at the Software Engineering Institute (SEI) believe this approach will lead to a more defensible and streamlined software development strategy. Although adopting this approach requires program office staff to have a level of model-based competency, using a digital model to review newly proposed capability(ies) provides a simple means to (1) review the capability with stakeholder(s) to ensure that there is adequate validation built into the process and (2) develop a higher quality artifact for negotiating with contractors.

The approach outlined in this paper should have broad applicability to acquisition programs, whether on the SWP or not, although there are specific artifacts referenced that are required for programs on the SWP. The suggested model-based approach elaborates elicited user need statements to enhance their usefulness. The benefits of this approach are illustrated by a use case that demonstrates the use of a model-based method to assess potential implementations of a newly proposed capability. Such an assessment would add value (or



reduce risk) for the users (source of the capability), the developers (implementers of the capability), and the acquirers. This approach represents a potential best practice for connecting stakeholder needs with practical reality using a digital engineering model as a facilitation tool.

The paper starts with a review of SWP and CNS guidance. It follows with an abstract discussion of capability and capability evolution. The next section discusses Digital Engineering (DE) and introduces the concept of a Digital System Model (DSM). The marriage of capability evolution and DSM is described with a notional model-based capability lifecycle. The paper concludes with a use case that illustrates different examples of how an organization might implement a DSM to elaborate its CNS along with a pragmatic discussion of risks and issues.

## Software Acquisition Pathway Context

The SWP, established by the Department of Defense (DoD) in 2020 and documented in DoD Instruction (DoDI) 5000.87, is part of the Adaptive Acquisition Framework (AAF; Defense Acquisition University [DAU], n.d.). The AAF was created in response to findings identified in several high-level studies, policy, and strategy documents, including the following:

- 2019 Defense Innovation Board Software Acquisition and Practices (Defense Innovation Board, 2019)
- Fiscal Year 2018 National Defense Authorization Act (NDAA) Section 873 and 874 Agile pilot programs (U.S. Congress, 2017)

Primarily, the SWP was created to accelerate the delivery of software-related capability to the warfighter. The SWP creates a separate mechanism for developing software capability that is independent of the system development lifecycle. The SWP promotes modern development practices, such as Agile and DevSecOps.

There are some important things to consider, especially with respect to Agile. As will be described below, capability evolves incrementally. Agile development practices fit this paradigm very well, as the Agile increments (i.e., sprints) are typically very short and produce capabilities that are testable. Project stakeholders (e.g., users, testers, developers) review the product at the conclusion of a sprint and can assess what is needed (if anything) to enhance the capability.

A digital, model-based representation of a capability should provide the “building blocks” that illustrate what to implement for that capability. These building blocks represent work to be completed. The work descriptions reside in a “Product Backlog.” The easiest way to think about this backlog is that it is like an inventory of work to be done. When the stakeholders build the project scope for a sprint, they draw from the backlog. (See a list of stakeholders and their roles in the appendix.)

Agile also defines terms, such as *minimum viable product* (MVP), *minimum viable capability release* (MVCR), and *definition of done*, that can have a dramatic impact on the as-defined inventory. For example, if we have a coherent definition of a capability, and we can represent it as “minimum viable,” we have a representation of a version of the capability that will enable the product to perform. The capability may not be completely implemented (as the user wished initially), but it may be “good enough.” By defining capability in this way, we create a trade space, or flexibility in the way the Program Management Office (PMO) can manage the project.

Priorities may change that require changes in the product backlog. If the PMO has an implementation that is good enough, it can shift the effort originally allocated to this capability to something else that has a higher priority. So, some inventory may become obsolete or new



inventory may be added as a result of trading off priorities. These changes could be quite significant and illustrate how useful the Agile paradigm is in the context of building capability.

On March 6, 2025, the Secretary of Defense issued a memorandum regarding improving the acquisition process that stated, “To meet this challenge, I am directing all DoD Components to adopt the Software Acquisition Pathway (SWP) as the preferred pathway for all software development components of business and weapon system programs in the Department” (Hegseth, 2025).


A program on the SWP needs to develop the plan for building the system’s software. The plan will likely be informed by several things, including

- prior implementations
- a set of criteria for establishing development priorities
- a guiding document that describes the desired software capabilities of the system

The last item, also referred to as a CNS, is a required artifact defined by DODI 5000.87:

a high-level capture of mission deficiencies, or enhancements to existing operational capabilities, features, interoperability needs, legacy interfaces and other attributes that provides enough information to define various software solutions as they relate to the overall threat environment. (DoD, 2020)

DODI 5000.87 includes a graphic that illustrates the expected high-level process for a program on the SWP. Note that the capability needs are shown in the upper left corner, but in the lower right corner, it says, “Software is never done.” The implication is that software capability is never done either. For example, new use cases are imagined after initial capabilities are deployed, threats evolve, hardware is updated, etc. All of these represent potential updates to a CNS. Therefore, it is important to think about the CNS as a living document.



[Source: (Department of Defense, 2020) ©2020 by Department of Defense. Reprinted with permission.]

**Figure 1: The Software Acquisition Pathway**

## Capability Needs Statement

A CNS is a key acquisition artifact that documents the high-level capabilities provided by the software. It is a living document and should be reviewed and updated as needed (at least annually), typically after a system's *value assessment* (Software Engineering Institute, 2026). A value assessment in Agile is the process of measuring the tangible benefits—such as increased functionality, efficiency, or revenue—delivered by teams. Such an assessment is usually reviewed every 10 months (or more often) to ensure its relevance. It focuses on prioritizing working software and customer feedback over rigid plans, using assessments to improve delivery and secure funding (DAU, n.d.).

To enter the SWP, a program must have a draft CNS. This is needed to get an SWP Acquisition Decision Memorandum (ADM) from its Decision Authority (DA). The program's sponsor is responsible for creating the draft CNS with input from users. The draft CNS can be based on a Concept of Operations (CONOPS), a Concept of Employment (CONEMP), or information from a Capabilities Based Assessment or Analysis of Alternatives (Warfighting Acquisition University, n.d.). The CNS is updated as needs change, as the capability is implemented, or as priorities change.

The CNS should also be clear and concise. It should describe the desired mission outcomes and how the software capabilities will enable these outcomes, and it should be fewer than 10 pages long. For guidance on developing your CNS, refer to *Cracking the CNS Code* (Software Engineering Institute, 2026). (See the CNS template for more descriptions [DAU, 2021].)

## Capability Definition Is Not Static

When a new capability is provided to users, it rarely works exactly the way the user (who requested the capability) expects. That is not necessarily bad, and this situation can't be solved by policy. The fact is that it is very difficult to define exactly what is needed. Further, the more you use something, the better you get at using it. Capabilities naturally evolve; they should be developed incrementally and be validated by users as they are developed. In this way, useful capabilities can be acquired in a systematic way.

For most capabilities, this means that there ought to be a way to characterize a capability so that it is defined as a “base” set of features plus feature enhancements (that potentially increase the value of the capability). In this way, the user has a chance to evaluate the implementation of the capability while the system is still in development. It is a challenge to precisely define these enhancements a priori, because it is hard to envision how to best use something until you are using it. We have seen this repeatedly with many products.

The safest approach is to first develop what everyone agrees that the capability must perform. Once that has been accomplished, and it is potentially deployed into an evaluation environment, enhancements will be easier to identify, understand, and prioritize.








**Figure 2: Capability Abstraction**

Figure 2 illustrates two ways to represent a capability. Think of a capability as something that needs to evolve. The representation on the left might reflect the initial (user's) characterization of the capability; but to implement it, the capability needs to be broken down into what must be done initially (often referred to with the descriptor "minimum viable") and what enhancements might be added in the future to make the capability better.

The transformation from the blue circle on the left to the discrete picture on the right provides two very important functions. There is a reviewable artifact that shows how the capability might be satisfied and a notional allocation that identifies the core functions and the functions defined as enhancements. Validating this picture with the user ought to be illuminating. The user may not wish to compromise and will insist on an "all (i.e., core) or nothing" strategy, but our experience has been that users don't think about capability in this way. It may be that there is a give-and-take to define the minimum viable capability.

Another way to think about capability evolution is that the developer may be implementing only the core capability but is probably aware of the enhancements that are on the horizon. Accordingly, the developer may propose to scope the minimum viable capability as the core plus some additional infrastructure that would need to be present if/when an enhancement is prioritized. This would be a judgment call for the acquirer.

For example, a fighter jet may need a radar-sensing capability. There are plenty of radars to choose from, and the minimum viable capability should be straightforward to implement. However, it may be that threat analysis suggests enhancements to improve survivability. For a radar system, these enhancements might relate to things such as integrating a jamming capability or increasing the resolution of the sensor to enable the detection of Small Unmanned Aerial Vehicles (SUAVs). Others may emerge during the acquisition that require a change in priority or an entirely new scope. As the capability evolves, enhancements are incorporated into the base as illustrated in Figure 3.




**Figure 3: The Evolution of Capability**

Based on the plan to incorporate Enhancement #3, the developer may have included some infrastructure to make the enhancement easier to implement. The enhanced capability, represented by the inner light blue circle, is slightly larger than the previous minimum viable capability, as Enhancement #3 has been incorporated into the plan. At any point during system (or software) development, stakeholders may declare that the capability is “done” primarily because it is good enough and other capabilities have a higher priority. This way of rationalizing how capability evolves fits right in with the Agile concept of MVP.

There is a history in DoD acquisition of missing the mark on capability implementation. When compromise isn’t possible on defining the core, there is no other choice than to build the capability in its entirety. What usually results is a capability implementation that looks like **Figure 4**. The capability is implemented with Enhancement #2 and Enhancement #3 incorporated, but Enhancement #1 and Enhancement #4 were not. This might have been caused by running out of time (or money) or some system dependency that was not implemented. Regardless, instead of developing a minimum viable capability, we are left with a partially implemented capability that potentially limits end-to-end system performance, which is not what we want to do.




**Figure 4: Undesirable Capability Implementation**

## Digital System Model

Digital Engineering (DE) is becoming pervasive as part of the development of software systems. Initial guidance for DE was issued by the DoD in 2018 and updated in late 2023 as DoDI 5000.97 (DoD, 2023). A handbook titled *Digital Engineering and Modeling Practices* (MIL-HDBK-539A) was just released in March 2026 (Department of War, 2026).

The basic value proposition for a DSM is to enable a process to shift left. Shifting left can be rationalized generally as a trade-off: early investment in a model (shifting left of cost) to prevent late identification of defects or issues. The economic case for DE is based on early defect detection, which was argued by Feiler et al. (2013) and Hansson et al. (2018). In short, these researchers present evidence that, in the domain of embedded safety-critical systems, 35% of errors are introduced in requirements, and a further 35% of errors are injected in architectural design. The cost of correcting an issue in later phases can be 300 to 1,000 times the cost of correcting it in phase (Dabney, 2003).

An organization that wishes to apply DE principles to its process seeks to provide early feedback that guides future development to mitigate project risk. These risks typically are manifested as cost and/or schedule overruns coupled with potential compromises in delivered capability.

A DSM that employs a systems modeling language (SysML) can document requirements and requirement elaborations. These elaborations typically take the form of block diagrams, sequence diagrams, and interface descriptions. An important element of a SysML model is that it maintains traceability between the diagrams and the requirements. This is very important

when reviewing new requirements or changes to existing ones. In the context of this paper, the DSM we refer to is one that is created and maintained by the acquirer.

The representation of a CNS in a DSM, created by the government, would be a break from tradition. We refer to this representation as a Government Reference Architecture (GRA). The GRA enables the PMO to elaborate capabilities and/or requirements in a model-based environment. To be clear, for this type of model application, the program office is assumed to represent the requirements with more than text. In other words, the acquirer is expected to elaborate on how it envisions the requirement to be implemented. The elaborated model-based representation can then be used for different purposes (e.g., requirement validation [with the requirement provider], requirement implementation guidance [for the project team]).

### Model Development Lifecycle

One of the most important things to do when attempting to develop a useful model is to proactively identify the way the model is intended to be used by its stakeholders. The reason for this is that a model development effort can go “off the rails” (generally because the model tool is so flexible) if model developers start to use the model in ways that it was not intended for. Model development represents a process, and there will likely be a trade-off in development between the normal process improvement activity of process review and improving and changing the intended use of the model. For example, the process review could identify that the model would have been more useful if a certain type of information was added. This may be true, but it would come at a cost (more time, more money, training, etc.). The model team should be very careful about changing the use cases of the model. A prototyping effort might be conducted very early in the project to validate the approach.

If model use case(s) are defined up front, then the process practitioners will have a structure they can work with to improve the way they use the model to perform their job. Changing the use cases after significant investment has been made in the model is generally not a good idea. So, we strongly advocate that the PMO should identify very early how it intends to use the model. Likewise, the PMO should communicate model use case(s) with others who will have a role in the process.

Beyond these tasks, the model developer should anticipate inputs and develop the process for (1) updating the model, (2) verifying that model updates reflect new inputs, and (3) releasing the updated model to stakeholders.




**Figure 5: GRA Model Development Overview**

A very simple representation of the process is depicted in Figure 5. Initially, the process flows from left to right, as the acquirer represents and elaborates the stakeholder’s capability needs in the GRA model. However, as the product is built, the process will flow in the other direction, reflecting design decisions and constraints that may impact the implementation of some of those needs. In this way, the user, the acquirer, and the contractor all participate in a verification and validation process that maintains consistency between user needs and the development of the system.

## What the Model Could Be Used to Do

The following are the most likely uses of the GRA:

**Validate User Needs.** The PMO will use its domain knowledge to build a notional model-based representation of each user need. This may involve anticipating specific hardware, software components, and interfaces. As part of the validation, the PMO attempts to characterize the relative priority of each user need. Documenting specific constraints (e.g., performance, cost, schedule, technology dependencies) should also be incorporated.

**Communicate User Needs to the Contractor.** The PMO incorporates the GRA into the contract it negotiates with its development contractor.

**Serve as a Communication Mechanism.** The PMO uses the GRA to communicate with users and the project team as development progress is made and priorities change.

**Provide Traceability Between User Needs and Project Backlog.** The PMO translates from the CNS into the GRA, identifying elements that need to be implemented. The work associated with building the capability traces from the GRA to the Project Backlog tool.

## Discussion of Capability Evolution and a Capability Lifecycle

Sometimes capabilities are explicit, and all stakeholders understand the capability. Sometimes the capability is something new. Let's examine the capability of an automobile's cruise control. At this point, everyone understands the concept behind the capability. The driver can maintain a constant speed without using the accelerator or the brake. The concept is very old; the first patent for a "Speedostat" was granted in 1950 to Ralph Teeter, and it basically accomplished the task of controlling the speed of the car (Sears, 2018). Cruise control was first released in 1958 by Chrysler. By the mid-1970s, cars with cruise control were being sold by American Motors, General Motors, Chrysler, and Ford.

The core capability in this case is to automatically maintain a constant speed of a car. Implementing the capability was initially accomplished in several different ways. A breakthrough occurred in the late-1980s when Motorola developed an electronic cruise control unit referred to as the Automotive Speed Control Processor. When integrated with an electronic throttle control, this unit became the starting point for modern cruise control, as it integrates the cruise control function with the engine management systems.

For several years after the release of the electronic unit, this capability remained fairly static. However, in the past 20 to 30 years, there has been a steady evolution of new capability enhancements. For the most part, these enhancements were enabled by new technology, but some were simple software enhancements. An example of a software enhancement is the way the cruising speed is adjusted, whether by 1 mph or by 5 mph, with each press of the up/down arrow.

When new technology (e.g., GPS, low-cost radar, lidar, or imaging systems) became available, new cruise control capabilities could be imagined. It is now common for cars to have *adaptive* cruise control. This capability automatically adjusts the speed of the vehicle to the one in front of it to maintain a safe driving distance. It is also normal now for the vehicle to have some sort of safety-oriented capability that the new cruise control capability integrates with.

New technology will continue to drive the evolution of this capability. An example of this might be a predictive adaptive cruise control system that gathers intelligence about the road ahead and adjusts the speed to deal with the specific driving context (e.g., reduce speed if a policeman is reported ahead or if weather conditions ahead may cause an unsafe driving condition).



For the Department of War (DoW), we should think of capabilities in these terms. There is a base capability: the ability to maintain the set point of the car's speed. Then we can think of enhancements to meet various purposes (e.g., accuracy, user interaction, new technology incorporation). The timeline for capability development might not take as long as it has for cruise control. In fact, new enhancements might be enabled within a single increment of development. This fits in perfectly with the Agile project management construct described earlier. However, thinking about capability in this way is different from the way it has been thought of in the past.

### **Model-Based Capability Definition**

For a specific acquisition, we can imagine a software capability that is required for the system. In most cases, this will not be an entirely new capability, rather it will be some legacy capability that needs to be enhanced. Enhancements are typically documented as requirements. These requirements come from users or, as in the case of security and/or safety, from a regulatory agency or the desire for compliance with DoW policies.

The main point of using interconnected DSMs is to ensure that mission needs, requirements, design, and verification are end-to-end traceable and analyzable using a structured modeling language rather than disconnected documents. A notional workflow for the capability translation might include the following steps: analyze mission threads in a model, identify gaps and perform trade studies, define new capabilities, analyze key performance characteristics and measures, allocate new capabilities to systems and identify impacted interfaces, and construct a high-level verification strategy.

Most systems being developed for the DoW are not greenfield (i.e., entirely new), although some system capabilities may be. What this means is that, in the most common case, the PMO has a very good picture of what it is acquiring.

For example, the Army has attempted to replace the Bradley Fighting Vehicle (BFV) several times over the last 25 years. Each time it tried, a PMO was stood up to manage the acquisition. The BFV replacement (especially from a software perspective) will have a similar capability as the existing BFV. There will be new technology, advanced computing capability, improved network communications, etc., but it should be easy for the acquirer to logically organize the BFV replacement vehicle based on the current BFV. For example, consider the navigation capability. It may be that the BFV did not incorporate GPS into its driver navigation, but a BFV replacement would require it. Regardless, the core logical navigation functions (e.g., waypoint navigation, reporting position) should be very similar.

If we can imagine how the PMO staff could logically construct the vehicle based on functional capabilities, it does not seem to be a stretch for them to use a model-based approach to perform the logical construction. The model may or may not act as a constraint for the project team, depending on how much control the PMO wants to exercise.

An example where the acquirer might constrain the contractor could be how to represent the BFV replacement's underlying infrastructure to comply with a modularity requirement. The PMO may want to employ modular components with a common infrastructure to simplify upgrading technology and other future modernization efforts. This decision will constrain the project team's ability to define modularity boundaries, along with the underlying infrastructure, and seek PMO approval as part of its process. While this may seem to be a completely normal part of the design process, there may be additional work performed by the project team to meet this requirement, and it is likely that the project team may not wish to comply.

In general, the capabilities can be represented in black boxes, with little to no constraint on what occurs within those boxes. Compliance with a data model may be required, and the boundary will be managed by an interface description. However, the PMO will focus on defining





the performance characteristics of the capability. In most cases, the project team is responsible for the precise details that describe how the black box capability is implemented.

The following sections define a notional process for how the discrete capability needs (as summarized in a CNS and defined by project stakeholders) could be incorporated into a model and used as a critical point of communication between the project team and the user community.

### Model-Based Capability Notional Process

Figure 6 illustrates a possible process for model-based capability. It's not a coincidence that operational needs are at the center of this diagram. Operational needs typically result from using the capability, observing some aspect of the capability that fell short, and recommending an enhancement to prevent a future recurrence of this deficiency. For the warfighter, this may have come from performing missions. A platform could have been used to perform some mission-related function with a negative result. An after-action review might have identified the problem and initiated the creation of a new need (or enhancement).

The pace of technology advancements seems to increase every day. This has always been a problem for DoD weapons systems, as these systems inevitably are outdated on the day they are released to the warfighter. One way to mitigate this risk is to actively elicit new operational requirements during the development of a new or updated system. This elicitation mitigates the risk of fielding obsolete products but increases the likelihood that the program will exceed its schedule and/or budget. It seems there will be some learning required to accomplish this process element effectively.




**Figure 6: Model-Based CNS Notional Process Overview**

For operational needs to make their way into the development pipeline, they need to be communicated to the PMO, and they need to be advocated by the sponsor. This activity implies that there is active collaboration between these process stakeholders. It also implies that there

are trade-offs that can be made *during an acquisition* to accomplish the new operational need without compromising the overall system capability.

To accomplish these trade-offs, the Agile notion of minimum viable capability must be embraced by all collaborators. In the best of worlds, a system is procured within its cost and schedule constraints and meets all of its functional requirements. When new requirements are actively elicited, there must either be a means of increasing the scope of the acquisition, or the new requirements need to be traded off against the existing project scope.

### **Capability Abstraction**

Capabilities are derived from stakeholder needs and refined into requirements. Some capabilities are so fundamental or widely accepted that they exist as baseline expectations even if not explicitly specified. Regardless, as operational needs are identified, an allocation is required to determine how the requirement will be satisfied. This allocation is an iterative and recursive problem, and it gets more complicated as the system matures. For example, the proposed implementation of a specific requirement may add memory or computational load to an already overallocated component, which might cause a new task or require refactoring the component to resolve the overallocation. It's important to know as early as possible what the implementation plan is and what options exist in the trade space.

The accumulation of requirements into capabilities results in the CNS. As stated earlier, the CNS is a living document that is required for a program to be on the SWP. We at the SEI strongly recommend creating the CNS collaboratively by the team members shown in **Figure 6**.

### **Model-Based Capability Translation**

For simplicity, let's assume that a GRA already exists for the system being acquired. As new operational needs are identified, the PMO can update the GRA, suggesting ways that a new need can be satisfied. Throughout this process, the PMO identifies risks early in the lifecycle, including suggesting possible mitigations (e.g., prototyping, a capability trade-off analysis, alternate implementations).

In principle, the GRA represents all of the system's capabilities and provides bidirectional traceability between the requirements and the specific elements of the GRA (e.g., block diagrams, behavior diagrams, interfaces, sequence diagrams). Although this representation is much more detailed than the CNS, it is important to remember that the linkage between the GRA and the CNS is explicit. The design choices or trade-offs that are in the GRA should be specifically traceable to the CNS.

### **Capability Definition Validation**

The proposed implementation of the operational need, represented in the GRA, is reviewed by the government team. The user validates that the proposed implementation is acceptable or identifies possible gaps in the implementation. The sponsor validates the implementation and can suggest other solutions from other projects it sponsors.

### **Updating the Roadmap and Product Backlog**

As updates are made to the CNS and priorities change, the PMO must update the Roadmap and the Product Backlog. These are artifacts of the Agile project management process, and they should accurately reflect the program's priorities and expectations. The Roadmap is a strategic artifact that represents big targets for the project team, and it is much more static than the Product Backlog. The Product Backlog should reflect tactics and low-level implementation strategies.



## Model-Based Acquisition

The PMO provides the GRA to prospective bidders for the system. This is probably done initially through a Request for Information (RFI), but it eventually becomes one of the Statement of Work (SOW) attachments. There are nuances that must be addressed to use a DE model that results in a successful acquisition. For example, the PMO must decide whether the GRA will continue to be updated after the contract is awarded (i.e., the Authoritative Source may transition to the Program Specific Model). The PMO might also require the project team to use the model-based contract attachment as a contract deliverable (the Contract Data Requirements List [CDRL]).

## GRA to Program-Specific Model Synchronization

Assuming that the PMO decides to maintain its GRA as the project team iterates its own system model, the PMO will need to update the GRA over the course of the acquisition to reflect the implementation decisions made and constraints identified (and approved by the PMO) by the project team. This work may require revalidating the capability with other project stakeholders, such as users.

The project team's system architecture will likely contain much more detail than the GRA, so some effort and competency will be required to accurately perform the synchronization.

## Project Product Validation

In the lower right portion of *Figure 6*, there is a representation of the delivery of the product into the government System (or Software) Integration Lab (SIL). This delivery is primarily shown for completeness, since programs on the SWP are required to deliver functional versions into the operational environment (or SIL) at least annually. However, project stakeholders can also evaluate the product in the SIL. In this way, they can provide the ultimate validation of the capability or identify gaps that need to be resolved to make the capability acceptable.

## Use Case – New Capability

The following use case illustrates the way a digital-engineering-based CNS can be integrated into the normal operations of an acquisition program. As you read the use case and the scenarios, think about the criteria you would use to assess the different implementations. For now, let's assume the following criteria:

- elapsed time for the capability to get to the warfighter (Faster = Higher)
- reliance on PMO technical competency (More Reliance = Lower)
- programmatic risk (i.e., cost, schedule, technical; More Risk = Lower)
- quality of the technical solution (Higher Quality = Higher)

## Use Case Context: New Capability

The user community for infantry fighting vehicles has observed an increase in the number of vehicles that have been damaged by undetected UAVs and needs to deploy a capability that protects the vehicle against this emerging threat. This community communicated this need to the sponsor of one of the PMOs for the ground vehicles, and the sponsor updated the CNS for the vehicle. The CNS update was provided on the normal update cycle. There were no other significant changes in the updated CNS.

The PMO has established a roadmap for the software effort with biannual releases of the platform software into its systems integration lab. The software contractor has adopted an Agile



software development lifecycle with sprints every 2 weeks. The PMO was already aware of this threat and had begun investigating alternatives for SUAV detection on the vehicle.

For this capability upgrade, the PMO assigns a team to investigate the alternatives for detecting SUAVs. The identified options include a modification to a legacy piece of equipment, a new piece of equipment, and a combination of a new piece of equipment and a modification to a legacy piece of equipment.

### **Scenario 1: Organic Model-Based Engineering**

The PMO team established to investigate alternatives for detecting SUAVs includes representation from Lethality, Systems, Software, and DE. The PMO invested in a GRA representation of the vehicle at the outset of the program. The team members were previously involved with the DE group when they initially built the GRA, so they are familiar with the model.

The team updates its model-based vehicle representation to build out various implementation options to support the capability. The options include the following four:

- Option 1: adding a new sensor expressly for the purpose of identifying SUAVs combined with existing lethality options
- Option 2: repurposing an existing sensor to use with existing lethality options
- Option 3: the same as Option 1 for sensing but with a new lethality option
- Option 4: the same as Option 2 for sensing but with a new lethality option

All of these options involve significant modifications to existing workflows on the vehicle, including target prioritization and target tracking. These workflows are elaborated by the team in the GRA, so the modifications can be represented as “red lines” over the existing content. The team assesses the modified workflows for each option using the information from the GRA and builds a table of potential impacts to the existing system. Beyond the size, weight, and power impacts, the team identifies that there may be impacts to the computing infrastructure, as all of the options require a machine learning algorithm with high-speed image processing. The team also identifies that significant impacts to the vehicle’s Tactics, Techniques, and Procedures (TTPs) will be required for some of the options, especially if repurposing existing equipment.

The team updates their assessment of the implementation options to include programmatic information (i.e., identification of cost and schedule budgets and project risks). It brings the GRA and its assessment to the stakeholders, reviews the potential implementations, and gathers information on how to define and develop the GRA representation of the core capability (as opposed to the enhancements). Throughout this process, the team observes that rapidly fielding an initial capability is the highest priority, so it focuses on the implementations that use existing hardware.

While reviewing the impact of the SUAV threat with the user from a mission performance perspective, it becomes clear that if the vehicle were able to share target data with its neighboring vehicles, it would significantly improve its ability to deal with this threat. This peer-to-peer communications capability does not currently exist, so it could not be included in the option set, but it represents a very desirable future enhancement to the counter-SUAV capability. The team integrates the new communications capability into the GRA as an enhancement.

The team works with its contracts team to initiate an RFI to identify possible candidates to implement this new capability and passes the high-level DSM to the candidates to help them decide whether to compete. The team elicits feedback from the candidates, which will be provided as part of its response.



The procurement and implementation will proceed normally, although it is highly likely that the winning offeror will be the one most familiar with the platform and its architecture, as selection of the technical solution will be weighted towards solutions that use existing hardware.

## **Scenario 2: Using an External Software Architect**

This scenario is functionally identical to Scenario 1 (above). The main difference is that the PMO uses an external subject matter expert (SME), a software architect with ground vehicle expertise, to manage the software architecture and update the GRA as required. This is a critical distinction between these two scenarios. Both DE tool-related skills and appropriate software domain expertise are required of the software architect to perform successfully in this role. It's entirely possible that the PMO could outsource the architecture work at the start of the acquisition and aim to transition this role to organic PMO staff as the project becomes more stable.

It's also possible that the sponsor organization (e.g., the Principal Executive Office [PEO]) may be looking to develop an architecture position for its portfolio of programs and may provide an incentive to the PMO for this purpose.

Regardless, the PMO team working on the new counter-SUAV capability develops and investigates the alternatives, working with the external software architect, and proceeds in much the same manner as what is described in Scenario 1.

## **Scenario 3: Software Architect by Project Team**

In this scenario, the PMO relies on the project team to develop the alternatives. This scenario is viable, but it places a burden on the PMO to ensure the identified alternatives are not overly dependent on solutions that favor the project team's technology. This is a natural conflict of interest (COI) for the project team.

The PMO needs to establish a consulting contract with the project team to support the PMO's need for changes to priorities and new requirements. In this way, the PMO can avoid this likely COI, as the project team could staff this consulting contract with a team that is firewalled.

## **Scenario Discussion**

These three scenarios are all valid ways for a PMO to implement a DSM to improve the quality of a CNS. The first two implement a GRA that is maintained by the PMO. The third is a contractor-maintained architecture, which is potentially the same as the system architecture but at a higher level of abstraction.

The first scenario relies on the competency of the PMO staff to identify the technical solution(s) and to provide the DE skills to develop the GRA. Assuming these competencies were properly provided, this approach should get a higher quality technical solution to the warfighter quicker. However, there is some risk, as the project team will need to cooperate.

The second scenario is less reliant on PMO competency but requires a contract action (to get the external software architect tasked to support this). This might add a bit more time when compared to the first scenario. To a certain extent, using an external architect is like what a person might do when implementing a home improvement project. If you don't want to design the improvement yourself, or don't trust your contractor to do it, you use a third party that is knowledgeable to provide an independent perspective.

The third scenario relies on the project team to provide the independent high-level design. This would also require a contract action of some kind to get the project team working on the new capability. The potential impact to the overall project might be higher in this case, as there is the possibility that the firewalled design team might not be as firewalled as we expect,





which increases the likelihood of a COI. There is also a predisposition on the part of a contractor-oriented project team to not be as open about their design solutions as a government-oriented one. There probably are several reasons for this, but this author's opinion is that it is mostly cultural.

## Summary

It's a big decision for a program to move its software development onto the SWP. It's an even bigger decision for the program to adopt a DE process to validate its CNS. This section of the paper summarizes the key considerations for the PMO.

The most important aspect of this approach is to accelerate the delivery of capability to the warfighter. We suggest that this is accomplished by maintaining a higher quality and validated representation of the CNS. On the surface, this sounds straightforward. However, the "devil is in the details." This approach adds burden to the PMO, and the process has lots of moving parts. As the system matures, dependencies emerge that need to be managed. If done correctly, the validation process will uncover capability-related risks and issues. When these risks and issues are identified early (e.g., before contract award), the project team has the best chance to mitigate them or make trade-offs with other capabilities.

To effectively use a DSM for this purpose, the PMO needs to keep some basic principles in mind. The PMO must be able to represent and communicate new capabilities within the DSM. This requires some model-based competency with a tool (e.g., Cameo), some specific software domain experience (for the particular software being developed), and the ability to translate the capability implementation alternatives from SysML into a form that the users understand. DE model use by PMOs has been increasing recently, typically because the PMO has required a DSM as a contract deliverable to define aspects of the product design (e.g., requirements, product architecture, component definitions, interfaces). As a result of this need to review the DSM, using DE model tools has become a necessary competency for some PMO staff.

One other point to consider is the amount of information in the DE model. As the DE model evolves, especially when it is under the control of the project team, a lot of details will be added. Most of these details are not necessary to understand the function of the software; they are needed to precisely specify the implementation. This low-level specification causes the model to be less useful for the purpose of interacting with users. The users need simpler (although functionally accurate) ways to envision end-to-end software threads. Keeping the representation simple helps the PMO explain alternative implementations in a way that users understand and can help to elicit critical aspects of the capabilities that may have been missed in non-model representations. The user-oriented model might be a different artifact than the project team model.

This last point raises an interesting conundrum. The implication is that there are two models being used to represent the same thing. This approach violates the DE Authoritative Source of Truth (ASoT) principle, which seeks to advocate a single source of truth as a DE goal. In the case discussed in this paper, the models are being used for different purposes. However, both models need to accurately represent the evolving system. This may be viewed as redundant by the PMO and other stakeholders who might ask, "Why am I paying for two models of the same system?"

No doubt there will be bumps along the way, but we need to remember that the reason why we are recommending this is to get software capability to the warfighter quicker. The enabling process is to improve the quality of capability definitions, so that we have a better understanding of minimum viable capability and can make informed decisions about trade-offs. This is accomplished by employing MBSE principles that transform the CNS from a static



description into an authoritative, model-based view of mission deficiencies, required and optional enhancements, interfaces, and interoperability. This view enables sponsors and users to prioritize and validate their needs iteratively with traceable evidence rather than with documents.

## **Appendix: Relevant Stakeholders**

The following are relevant stakeholders to be considered as part of the processes described in this paper (Software Engineering Institute, 2025).

### **Decision Authority**

The DA is the designated individual who has overall responsibility for a program. This individual has the authority to approve a program's progression through the acquisition process and is responsible for reporting cost, schedule, and performance results. For programs on the SWP, the DA documents the rationale and approves the decision to go on the SWP. In general, the DA approves strategies, reviews progress, and provides go/no-go decisions. In the context of this paper, the DA approves the CNS and reviews/approves the PMO's roadmap.

### **Program Sponsor**

The sponsor is the individual who, or organization that, holds the authority and advocates for needed end user capabilities and associated resource commitments. The sponsor should be part of a collaborative team, made up of the sponsor, users, and PMO personnel, that jointly elaborates operational needs and develops (or updates) the CNS.

### **Users, Project Stakeholders**

The users and stakeholders are individuals who ultimately use the software. A user (or end user) conveys operational concepts, requirements, and needs; participates in capability prioritization and sprint (or increment) demonstrations; and provides feedback on developed capabilities. The users should be part of a collaborative team made up of the sponsor, users, and PMO personnel jointly developing the CNS.

### **Program Management Office**

The PMO should be part of a collaborative team made up of the sponsor, users, and PMO personnel jointly developing the CNS. The PMO needs to ensure the CNS is written at a high enough level to allow trade space during Agile software development. The PMO also ensures that relevant areas (e.g., logistics, training, cybersecurity, certifications) are included in the CNS. The PMO, along with the sponsor, must ensure that any changes in needed capabilities or threats trigger an update to the CNS.

### **Contractor(s) or Project Team**

A multidisciplinary team is typically funded to develop a system, software, or other warfighting capability through a formal mechanism (e.g., contract, agreement, interdepartmental request). This team can be a government team or a contractor/vendor team. It includes roles such as project manager, development team member, system integrator, finance specialist, contracting specialist, and product test and evaluation team member. Depending on the program, this team also might include a product owner who shares that role with or interacts with the product owner in the PMO.

[DISTRIBUTION STATEMENT A] Approved for public release and unlimited distribution.

<https://www.sei.cmu.edu>

Carnegie Mellon University 2026

This material is based upon work supported by the Department of War under Air Force Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering



Institute, a federally funded research and development center.  
 NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN “AS-IS” BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.  
 [DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution.  
 Carnegie Mellon® is registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.  
 DM26-0252

## References

- Dabney, J. B. (2003). *Return on investment of independent verification and validation study* (Preliminary Phase 2B Report). NASA.
- Defense Acquisition University. (n.d.). *Adaptive acquisition framework*. Retrieved March 11, 2026, from <https://aaf.dau.edu/>
- Defense Acquisition University. (n.d.). *Define capability needs*. Retrieved March 11, 2026, from <https://aaf.dau.edu/aaf/software/define-capability-needs/>
- Defense Acquisition University. (2021, May 14). *Capabilities Needs Statement (CNS) template*. <https://aaf.dau.edu/storage/2022/06/CNS-Template-14-May-2021.docx>
- Defense Innovation Board. (2019). *Software Acquisition and Practices (SWAP) study*. DoD. [https://media.defense.gov/2019/Mar/07/2002097482/-1/-1/0/SWAP\\_STUDY\\_VIGNETTES.PDF](https://media.defense.gov/2019/Mar/07/2002097482/-1/-1/0/SWAP_STUDY_VIGNETTES.PDF)
- DoD. (2020). *Operation of the software acquisition pathway*. [https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500087p.PDF?ver=viRfQj4v\\_LgN1JxpB\\_dpA%3D%3D](https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500087p.PDF?ver=viRfQj4v_LgN1JxpB_dpA%3D%3D)
- DoD. (2023). *Digital engineering*. [https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500097p.PDF?ver=bEPlqKXaLUTK\\_lu5iTNREw%3d%3d](https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500097p.PDF?ver=bEPlqKXaLUTK_lu5iTNREw%3d%3d)
- DoW. (2026). *Digital engineering and modeling practices*. <https://quicksearch.dla.mil/Transient/8606E424EB0D44868358B4367CE254E1.pdf>
- Feiler, P. H., Goodenough, J. B., Gurfinkel, A., Weinstock, C. B., & Wrage, L. (2013). *Four pillars for improving the quality of safety-critical software-reliant systems*. Software Engineering Institute. <https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=47791>
- Hansson, J., Helton, S., & Feiler, P. H. (2018). *ROI analysis of the system architecture virtual integration initiative*. Software Engineering Institute. <https://www.doi.org/10.1184/R1/12363080.v1>
- Hegseth, P. (2025). *Directing modern software acquisition to maximize lethality*. <https://media.defense.gov/2025/Mar/07/2003662943/-1/-1/1/DIRECTING-MODERN-SOFTWARE-ACQUISITION-TO-MAXIMIZE-LETHALITY.PDF>
- Sears, D. (2018, March 8). The sightless visionary who invented cruise control. *Smithsonian Magazine*. <https://www.smithsonianmag.com/innovation/sightless-visionary-who-invented-cruise-control-180968418/>
- Software Engineering Institute. (2025). *Software acquisition go bag user's guide*. PA: Software Engineering Institute. [https://www.sei.cmu.edu/documents/6412/Go\\_Bag\\_Users\\_Guide.pdf](https://www.sei.cmu.edu/documents/6412/Go_Bag_Users_Guide.pdf)
- Software Engineering Institute. (2026). *Cracking the CNS code*. [https://www.sei.cmu.edu/documents/6465/Cracking\\_the\\_CNS\\_Code.pdf](https://www.sei.cmu.edu/documents/6465/Cracking_the_CNS_Code.pdf)



U.S. Congress. (2017, December 12). *National Defense Authorization Act for Fiscal Year 2018*. Office of the Law Revision Counsel. <https://www.congress.gov/bill/115th-congress/house-bill/2810>

Warfighting Acquisition University. (n.d.). *Analysis of Alternatives (AoA)*. Retrieved March 11, 2026, from <https://www.dau.edu/acquikipedia-article/analysis-alternatives-aoa-0>

Warfighting Acquisition University. (n.d.). *Capabilities-Based Assessment (CBA)*. Retrieved March 11, 2026, from <https://www.dau.edu/acquikipedia-article/capabilities-based-assessment-cba-0>



# Transforming Acquisition for Speed, Agility, and Effectiveness

**Brigid O’Hearn**—Principal Engineer at CMU’s Software Engineering Institute (SEI), brings 30 years of software experience to Department of War (DoW) software acquisition challenges. As the Software Modernization Policy Lead for the SEI, O’Hearn influences the DoW’s Software Modernization Strategy and Implementation, software workforce initiatives, and software/cyber policy. Her research leadership is informed by 20 years as a DoW civilian supporting the U.S. Air Force and U.S. Navy. She was the first chief software officer in the F-35 Joint Program Office and is a proud graduate of National Defense University’s Eisenhower School and DoW’s Defense Senior Leader Development program. [bpohearn@sei.cmu.edu]

**Carol Smith**—Principal Research Scientist at CMU’s Software Engineering Institute, leads research to improve the user experiences (UX) of the nation’s artificial intelligence (AI) systems, autonomous systems, and other emerging technologies. Smith has conducted research to improve interactions with complex systems across industries for over 25 years. She has been involved with developing trailblazing responsible approaches to AI and autonomous system development since 2015. Smith is a co-author of the DoW Defense Innovation Unit Responsible AI Guidelines and is recognized globally as an advocate for human-centered AI. Smith holds an MS in Human-Computer Interaction from DePaul University. [cjsmith@sei.cmu.edu]

**Rita Creel**—Principal Engineer at CMU’s Software Engineering Institute, works with programs adopting the Software Acquisition Pathway to manage risk and amplify effective practices. Her 35 years of experience in software-intensive systems engineering, acquisition, and cybersecurity spans technical, management, and organizational leadership roles. In industry, she led development of software for on-orbit satellite systems, anomaly resolution, and integrated satellite systems factory test. Creel’s current research interests include complexity leadership, goal-driven measurement, digital engineering, and multimodal machine learning. She has BS and MS degrees in computer science and an MA degree in organizational leadership. [rc@sei.cmu.edu]

## Abstract

Acquisition programs face several common challenges in adopting the Department of War (DoW) Software Acquisition Pathway (SWP) and acquiring artificial intelligence (AI) capabilities. This paper describes these challenges and the work underway to address them and help programs move swiftly from high-level policy to battlefield-ready practice.

## Introduction

Over the past year, the urgent need for acquisition reform has been a strong and recurrent theme in White House, Department of War (DoW), and Congressional communications. The April 9, 2025, White House Executive Order, *Modernizing Defense Acquisitions and Spurring Innovation in the Defense Industrial Base* (The White House, 2025), laid the groundwork for the Secretary of War’s (SecWar’s) *Acquisition Transformation Strategy* (DoW, 2025) and the associated memo, released on November 7 (Hegseth, 2025b). These documents include phrases such as “accelerate defense procurement,” “emphasis on speed, flexibility, and execution,” and “state-of-the-art capabilities at speed and scale” (The White House, 2025, p. 1).

Prior to the release of the memos cited above, the SecWar’s March 6 memo, *Directing Modern Software Acquisition to Maximize Lethality* (Hegseth, 2025a), emphasized software’s critical role “at the core of every weapon and supporting system we field” (p. 1) and the need to immediately shift to a modern software acquisition construct. Leveraging advanced technology and the commercial ecosystem, modern software acquisition practices provide the ability to deliver, scale, and evolve capabilities to outpace our adversaries.





This focus on modernizing and accelerating software acquisition aligns with Senator Roger Wicker's report, *Restoring Freedom's Forge: American Innovation Unleashed* (Wicker, 2024) and is reflected in the Senate's *Fostering Reform and Government Efficiency in Defense (FoRGED) Act* (2024) and in the House of Representatives' *Streamlining Procurement for Effective Execution and Delivery (SPEED) and National Defense Authorization Act for Fiscal Year 2026* (2025). Five acquisition transformation themes are woven through these memoranda, policy changes, and proposals:

1. Eliminate bottlenecks and barriers.
2. Encourage high levels of innovation.
3. Support rapid decision-making and intelligent risk-taking.
4. Revitalize the defense industrial base.
5. Streamline budgeting/contracting processes.

This paper reports on research findings and recommendations that align with the above themes to promote speed, agility, and effectiveness. It begins by discussing the genesis of the Software Acquisition Pathway (SWP), which operationalizes the five themes for software. Next, the paper describes common SWP adoption challenges that DoW programs have encountered, and a collection of techniques and tools that can facilitate successful outcomes. In addition, with an expanding drive to adopt artificial intelligence (AI) enabled capabilities, the paper presents results from a workshop on AI acquisition, including effective practices for acquiring AI-enabled capabilities (Rishel et al., 2025).

## **Challenges in Software Acquisition and AI Adoption for National Security**

Modern software engineering approaches reflect the four fundamental values articulated in the *Agile Manifesto* (Beck et al., 2001): Interactions and people over processes and tools, working software over comprehensive documentation, customer collaboration over contract negotiation, and responding to change over following a plan. These values recognize that software's real power is its inherent flexibility, and its real value is in meeting user needs. However, until recently, DoW lacked a policy framework to realize the full benefit of this flexibility and its potential to deliver game-changing capabilities to the warfighter.

### **Software Acquisition Pathway (SWP): Capability and Agility at the Speed of Need**

Traditional DoW acquisition frameworks were designed for hardware, not rapidly evolving, dynamic software. Aligned with the manufacturing- and production-oriented needs and timelines of hardware, these frameworks revolved around infrequent, major milestone reviews rather than rapid, iterative feedback loops. Interactions with user representatives were not codified, value was not measured, and neither was a requirement to "demonstrate the viability and effectiveness of capabilities for operational use not later than 1 year" (DoD, 2020, p. 3) after funds are first obligated for development. Without early opportunities for user feedback or to work out integration issues, programs have struggled to deliver software capability at the speed required for modern missions.

Beginning with the Software Acquisition and Practices (SWAP) study, *Software Is Never Done: Refactoring the Acquisition Code for Competitive Advantage* (Defense Innovation Board, 2019), the DoW and SEI sought to develop policy constructs more suitable for software. They designed and implemented pilot programs used to inform, test, and refine the policy constructs underlying the SWP, which are defined in DoD Instruction 5000.87, *Operation of the Software Acquisition Pathway* (DoD, 2020). The SWP is one of several acquisition pathways introduced in DoD Instruction 5000.02, *Operation of the Adaptive Acquisition Framework (AAF)* (DoD, 2022),



a major revision of policy that establishes a foundation for adapting acquisition processes to a variety of contexts.

The objective of the SWP is to enable capability delivery to the warfighter “at the speed of mission need” (Defense Innovation Board, 2019, p. i), which eliminates unneeded and onerous requirements. Accordingly, the SecWar directed the adoption of the SWP “as the preferred pathway for all software development components of business and weapon system programs in the Department” (Hegseth, 2025a). The SWP emphasizes modern software practices focused on frequent, iterative delivery cycles that enable user feedback and update early in the life cycle, when cost and schedule impacts are much lower. In the end, the target result is a delivered capability that meets warfighter needs on schedule and within budget.

The SWP offers great potential benefits to the acquisition community. However, it represents a significant departure from previous software acquisition practice and requires a commitment to learning and practicing a new—and better—way of working. Based on work with programs adopting and learning about the SWP, the SEI identified five challenges commonly encountered along the way. These are described in Table 1.

**Table 1. Key Barriers to Adopting the Software Acquisition Pathway (SWP)**

| <b>SWP ADOPTION CHALLENGE</b>                                 | <b>DESCRIPTION</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SWP Readiness</b>                                          | Entering the <i>SWP Planning phase</i> requires a signed Acquisition Decision Memorandum (ADM) and draft Capability Needs Statement (CNS) (Warfighting Acquisition University, n.d.-a, n.d.-b, n.d.-c). But SWP success requires additional preparation for the program team to be adequately organized and equipped for successful execution. It’s necessary to gain an adequate understanding of what modern software development approaches and Agile methodologies look like in practice as well as how to provide a framework that enables them to function well in your program and facilitate early detection when execution drifts off course.                                                                                               |
| <b>Capability Needs Statement (CNS) (Requirements) Issues</b> | Many users and program office teams struggle to shift from a <i>Capability Development Document (CDD)</i> articulation of needs to a CNS at the appropriate level (Warfighting Acquisition University, n.d.-b, n.d.-c). In some cases, the result is a CNS that is merely CDD content that has been copied and pasted into the CNS template, making it overly detailed. The specificity required of a CDD runs counter to the flexibility and adaptability the SWP is designed to enable. In other cases, the CNS lacks the information needed for foundational system architecture work (e.g., reliability, testability, and maintainability requirements). An appropriate CNS provides the right level of detail to enable the needed flexibility. |
| <b>Poorly Specified User Agreements (UAs)</b>                 | A UA is a vital handshake that explicitly links the development of capabilities to real-world warfighter outcomes. UAs are closely linked to Value Assessments (VAs), which provide feedback and return on investment data on the value delivered. Often, UAs are implemented at the wrong level (i.e., too far removed from the day-to-day work of users), or they don’t sufficiently consider different types of end users who have a stake in the capability being delivered.                                                                                                                                                                                                                                                                     |
| <b>Confusion About Software and Program Metrics</b>           | Program teams often learn about a slip in a delivery date or a cost overrun too late to do anything about it. Metrics provided by product teams may fail to expose risks early enough to mitigate them before they become problems that are difficult and costly to resolve.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Ineffective Scoping of Capability Releases</b>             | Many program teams begin their SWP journey without an adequate understanding of Minimum Viable Product (MVP) and Minimum Viable Capability Release (MVCR). An inadequate definition and scoping of a release’s MVP and MVCR can lead to misaligned expectations between the program team and stakeholder groups (e.g., warfighters, certifiers, testers).                                                                                                                                                                                                                                                                                                                                                                                            |



Proven practices for overcoming the challenges identified in **Error! Reference source not found.**, along with application guidance, are provided in the “Overcoming the Challenges: Acquisition Resources and Effective Practices” section of this paper.

Since the SWP’s introduction in 2020, a burst of innovation and investment in AI produced results that turned DoW interest into an urgent pursuit of AI capabilities. The software acquisition process can accommodate AI-enabled software acquisition. However, AI is inherently different from conventional software in its development cycle and its behavioral characteristics. AI software development includes activities and artifacts that are not part of conventional software development. In addition, some conventional software life-cycle activities must be modified for AI. The next section discusses AI adoption challenges related to differences between AI and conventional software.

## Ubiquitous Artificial Intelligence (AI): Adoption Challenges for AI Acquisition

The field of AI research was founded in the 1950s with capabilities and technologies advancing in spurts (Stone et al., 2016, Appendix I). The disappointing pace of progress led to pauses in funding in the 1970s and 1990s; and then, almost 30 years ago, the pace picked up again. What is now called “traditional” AI primarily utilizes machine learning (ML) and has been in constant use since then, with a dramatic increase over the past 5 years in more advanced AI-enabled capabilities offered to the general public.

Large language models (LLMs) are an ML system that have quickly become ubiquitous. ChatGPT, released in November 2022, was the first publicly available LLM-enabled chatbot, and within 2 months it was estimated to have reached 100 million monthly active users (Hu, 2023). When Pew Research Center conducted a survey around that time, only 30% of U.S. adult respondents were able to correctly recognize six examples of AI in everyday life (Kennedy et al., 2023). In contrast, by 2025, another Pew Research Center survey showed that most American respondents (79%) had heard at least a little about ChatGPT, and 58% of adults under 30 had used it (Sidoti & McClain, 2025). However, ChatGPT and other LLM-enabled chatbots are only one of many types of AI people use each day, often without realizing it. Figure 1 depicts a hierarchy of AI categories and where ChatGPT fits.




**Figure 1. Artificial Intelligence (AI): Example Subfields, Approaches, and Architectures**

Like any tool or technology, some types of AI are better than others for certain tasks, and knowledge of these systems is necessary to effectively identify an AI technology suitable for a particular need. For example, LLMs are good at generative text-based tasks such as writing letters and summarizing content. However, they are not as good at retrieving facts or responding as consistently as other types of systems. LLMs respond with incorrect information more frequently than other systems and may combine unrelated pieces of information in unexpected and even nonsensical ways, which produce results often referred to as “hallucinations.”

### **Data Is the Cornerstone of AI-Enabled Systems**

A key difference between conventional software and AI-enabled systems is that for AI-enabled systems, data serves as the cornerstone, the foundation on which the algorithms, models, and architecture are constructed. To develop an ML capability, a data scientist or ML engineer provides a set of training data inputs and corresponding results<sup>1</sup> that have been reviewed by a subject matter expert, to an ML model. The model learns the patterns inherent in the data and the parameters are adjusted until a loss (error) function is minimized and correct results are produced. After the model has learned the patterns, the model can apply the learned patterns to produce correct results from new data (inputs it hasn't yet encountered).

Since AI-enabled systems are heavily dependent on data for their development and operation, identifying datasets that are relevant to a specific use case and at the appropriate classification is paramount for the system to be successful. For example, when there is a need to provide access to military doctrine, a helpful chatbot system may be the first choice. Chatbots can be fine-tuned and methods such as retrieval augmented generation (RAG) can further improve their responses. However, since LLMs, such as chatbots, are typically trained on the entire internet, they contain a lot of irrelevant information and will be slow compared to other options. Finally, preparing data for an AI system may require a considerable commitment of time and resources, and using AI may not be worth the time and cost. In this case, a different type of system may be much more effective.

### **Considerations for Using AI in National Security Contexts**

AI adoption for national security use cases is more challenging than for consumer products for several reasons. First, the overall risk is higher, particularly in terms of the more serious consequences of failure. A mistake in a commercial chatbot might cause momentary confusion. A mistake in an intelligence summary could lead to mission failure. Second, AI tools must often interoperate or be integrated with legacy systems, and these systems may not adequately support modern software. Finally, AI systems require relatively large sets of digitized data that have some variation and are relevant to the use case. Much of the data used in national security contexts is sensitive or classified and must be safeguarded at all phases of the AI life cycle.

These requirements can complicate development and maintenance of the model over time, particularly if the model is trained using sensitive or classified data. For example, changing conditions in the environment are likely to result in changes to the distribution of input data. As the input data diverges from the training data, the model's behavior will drift, and it will not produce correct results. To detect these kinds of changes, a monitoring system and human oversight are needed. To correct the situation, it may be necessary to fine-tune or retrain the

---

<sup>1</sup> When a machine learning model is given correct results along with training data inputs, it is using supervised learning. Known results “supervise” the model's predicted results and feed the error back through the model, which results in updates to the model's parameters. Other types of learning include unsupervised, semi-supervised, and reinforcement learning.





model. Having available staff and resources to collect data and do the fine-tuning or retraining in a timely manner, particularly if sensitive or classified, is likely to be challenging.

In addition to the above, Novak (2021) identified other acquisition- and policy-focused challenges in the literature. These challenges fall into a variety of categories and include technical, data, acquisition approach, adoption and culture, organizational, and legal and ethical challenges (pp. 2–3). Although the report is several years old, many are still relevant.

## Overcoming the Challenges: Resources and Effective Practices

Over the past year, the SEI has engaged with DoW acquisition organizations and programs and the software engineering community to develop and share practices and resources that address the challenges identified above. The Software Acquisition Go Bag (SEI, n.d.-b) is a resource that identifies essential, effective practices for navigating the SWP. These are described next, followed by a set of proven practices for AI acquisition that were discussed and developed during an SEI-led workshop (Rishel et al., 2025).

### Software Acquisition Go Bag: Delivering Software at the Speed of Need

The SEI has extensive knowledge and experience sharpening software acquisition practices based on decades of work with acquisition programs; most recently, it implemented the SWP in DoW programs and their partner organizations. In *Directing Modern Software Acquisition to Maximize Lethality*, SecWar Hegseth (2025a) amplified the DoW's commitment to moving programs to the next level of acquisition capability, a goal firmly aligned with the SEI's purpose. Determined to continue that mission and be instrumental in helping DoW programs and others improve their acquisition practices, in 2025, the SEI launched the Software Acquisition Go Bag—"Go Bag" for short (SEI, n.d.-b).

Go Bag is a system of resources designed to help users (e.g., program teams, user representatives, and their industry partners) confidently make acquisition decisions, including how to implement the SWP in their unique environments to achieve the desired outcomes. These resources include

- *Tactical Guides* that contain actionable information about a specific software acquisition topic
- *Live Webcasts and Podcasts* that enable participants to ask questions and share thoughts and are recorded for later viewing/listening
- *Blog Posts* that compile the latest news and Go Bag offerings
- *Templates and Other Tools* that support software acquisition activities

Go Bag enables users to collect and tailor the resources they need to devise strategies, practices, and execution plans to overcome their software acquisition challenges. You can "Pack Your Go Bag" with resources that include tips for following modern software acquisition, development, and deployment practices (SEI, n.d.-a).

To simplify SWP adoption, the SEI created an initial collection of resources called the SWP Essentials Kit (SEI, 2025). This kit includes fact sheets, blog posts, webcasts, tactical guides, supplements, and short video clips focused on concepts and artifacts that are essential to begin a successful SWP journey.

Table 2 describes the Go Bag resources aimed at responding to each SWP Adoption Challenge identified in Table 1. As of the date this paper was written, three of the resources listed in the table are published on the Software Engineering Institute Website





(<https://sei.cmu.edu>), another is in final publication, and two more are in progress and soon to be published.

**Table 2. Resources for Transcending Barriers to Software Pathway SWP Adoption**

| TO OVERCOME THIS CHALLENGE                          | TRY THESE RESOURCES                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SWP Readiness</b>                                | <p><i>What Is the SWP?</i> is an easy-to-follow guide to SWP basics.<br/> <a href="https://www.sei.cmu.edu/library/what-is-the-swp/">https://www.sei.cmu.edu/library/what-is-the-swp/</a></p> <p><i>SWP: Ready, Set, Go!</i>: The SWP allows for the rapid acquisition and delivery of software that provides warfighters/users with innovative capabilities they need to outpace our adversaries and win the next fight. Sounds like just what's needed, right? But it's true that the details of the SWP can get confusing. Since not all programs are a good fit for following the SWP, this guide helps you determine if your program fits with the SWP and if it is ready.<br/> <a href="https://www.sei.cmu.edu/library/swp-ready-set-go/">https://www.sei.cmu.edu/library/swp-ready-set-go/</a></p>  |
| <b>CNS (Requirements) Issues</b>                    | <p><i>Cracking the CNS Code: Getting the Capability Needs Statement (CNS) Right</i> unlocks substantial benefits, for example, with guidance on developing an acquisition strategy and cost estimates. Taking time to engage effectively with users and other stakeholders is critically important; skipping this or getting it wrong can create avoidable pitfalls, for example, making it more difficult to prioritize the backlog. User feedback on early prototypes through usability testing will help validate the CNS. This guide provides tips and examples for developing a CNS that addresses your program's unique objectives at the right level of detail.<br/> <a href="https://www.sei.cmu.edu/library/cracking-the-cns-code/">https://www.sei.cmu.edu/library/cracking-the-cns-code/</a></p> |
| <b>Poorly Specified User Agreements (UAs)</b>       | <p><i>In progress</i></p> <p>A UA is only as good as the solution that is generated. Cycles of iterative development and usability testing will provide confidence that the interface and data the system provides meet program needs and expectations. This tactical guide helps you develop an effective UA. Collaboration among the PMO, product team, and users is paramount to success in software acquisition, and the UA is the cornerstone of that collaboration.</p>                                                                                                                                                                                                                                                                                                                               |
| <b>Confusion About Software and Program Metrics</b> | <p><i>In final publication</i></p> <p>The <i>Pack Light, Measure Right</i> tactical guide and its supplements will support your program's ability to pack only the measures that "earn their keep."</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Ineffective Scoping of Capability Releases</b>   | <p><i>In progress</i></p> <p>This tactical guide will describe the intended content and purpose of three types of capability release: The <i>minimum viable product (MVP)</i>, <i>Minimum Viable Capability Release (MVCR)</i>, and subsequent <i>Releases</i>. This guide will also share the keys to successful <i>Release Planning</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

When the resources in the SWP Essentials Kit (SEI, 2025) are complete, the SEI will move forward to develop Go Bag resources on new topics as well as more detail, where



needed, on the topics in listed Table 2. One of the topics planned for a future Go Bag resource is acquiring AI-enabled capabilities and systems.

### Effective Practices and Guidelines Identified at the AI Acquisition Workshop

AI acquisition and implementation continue to be challenging for many programs. Common questions include why and when to adopt AI and how to choose an AI-enabled solution. In a 2025 workshop the SEI conducted, participants shared practices and guidelines they found useful in answering these questions, and below is a curated list (Rishel et al., 2025).

- **Start with a mission need.** Understand the problem to be solved, the context of use, and the end users' needs, and then choose a solution that addresses those needs and is likely to improve the situation. An AI-enabled solution may not be the best choice.
- **Select data.** Ensure datasets selected for training are appropriate and pass a robust assessment process, as described under "Ensuring AI Is a Match for the Operational Context."
- **Prepare data.** Data preparation can consume considerable time and resources in determining dataset quality and adequacy. In addition, missing values, imbalanced datasets, and outliers need to be handled. For some solutions, feature selection and engineering will be needed.
- **Manage data.** Of the many aspects of data management, security and protection of sensitive and classified information are particularly applicable to national security systems. Giving an AI solution access to a shared data repository with a variety of sources increases the risk of classified data spillage (Novak, 2021).
- **Ask how the model works.** Vendors should be expected to describe the training process of the model, the data it uses, how it makes decisions, and how and when it will be updated to a variety of audiences, both technical and non-technical. Disregard AI systems that are described as opaque or as a black box. Vendors do not need to divulge intellectual property to share this level of information.
- **Prioritize trustworthy systems.** AI systems are not reliable in the same way traditional software is. A trustworthy system provides an experience that matches end-users' expectations and is appropriately transparent to them. For example, an AI system should indicate to warfighters when its performance is degraded due to concept drift, or it is not working as intended. Trustworthy systems provide evidence for how outcomes are produced and meet (or exceed) performance metrics.
- **Choose modular systems.** Instead of seeking versatile solutions, identify tools that perform well on specific tasks, fit within the existing operating environment, and can be added or removed easily to reduce risk and prevent being tied to one vendor.
- **Run a pilot before issuing a contract.** Start with a small-scale experiment in the target environment with warfighters in realistic conditions. Partner with end-users to explore the system and to confirm that it does what they need it to. Continue to use the environment to evaluate the system prior to significant scaling in nominal and off nominal conditions, and throughout its deployment. Warfighters can also help refine requirements and identify risks, prior to updates of contract language.



- **Build in feedback and human oversight.** AI systems are dynamic by nature, and analysts and operators must be capable of identifying and reporting mistakes. Along with testing and evaluation efforts, AI systems need continuous monitoring for undesirable changes in model behavior (model drift) particularly in higher risk, sensitive, or classified environments. In addition, the system should notify an operator or administrator of issues so they can determine what action to take.
- **Plan for failure.** Even high-performing models will fail, and teams should understand the existing capabilities, constraints, and possible consequences. AI systems should be designed to be resilient so that they detect and recover from issues when possible, and the team should have a plan for when it cannot.

The above practices and guidelines suggest actions that acquirers can take to reduce risk and increase the likelihood that an AI-enabled system or capability will effectively perform its intended functions and be adaptable to change. While not comprehensive, they are broadly applicable and worthwhile areas for knowledge and skill building.

## Conclusion

The research described in this paper represents progress in two areas relevant to transforming acquisition for speed, agility, and effectiveness. The first is the Software Acquisition Go Bag, a collection of easy-to-use, accessible resources to facilitate software acquisition. The SWP Essentials Kit is the first set of resources in the collection, designed to accelerate and expand the successful adoption of the SWP, which enables deployment of adaptable software at the speed of need. Three of the five SWP Essentials Kit resources have been published on the Go Bag website, and additional resources are planned. New resources will be added to the website as they are published.

The exponential rise and sudden ubiquity of AI-enabled systems and capabilities have created an urgent need for guidance in acquiring AI. The SEI's AI Acquisition workshop produced a starter set of broadly applicable practices and guidelines for organizations to use, but more work is needed to produce a comprehensive set. As the practical use of AI expands and matures, acquisition and AI practitioners and researchers, industry associations, professional societies, and standards bodies have published and are continuing to develop comprehensive guidance documents and standards. Several of these adapt conventional development and acquisition phases, activities, and practices for AI governance, oversight, and procurement. For example, *IEEE Standard for the Procurement of Artificial Intelligence and Automated Decision Systems* (IEEE, 2025) focuses on procuring commercial AI products and services within a contractual framework, excluding development done in house or acquirer-supplier development.

With a plethora of approaches, use cases, contexts, and risks, organizations navigating AI acquisition must start with a solid understanding of the mission need and requirements of the operational environment, context, and constraints. In addition, they must have a solid grasp of the risks and uncertainties inherent in AI development and operational use. To speed the innovation and adoption of AI capabilities, blending this knowledge with the streamlined requirements of the SWP will provide a foundation for the rapid delivery of effective capability to the warfighter.

[DISTRIBUTION STATEMENT A] Approved for public release and unlimited distribution.

<https://www.sei.cmu.edu>

Copyright 2026 Carnegie Mellon University.



This material is based upon work funded and supported by the Department of War under Air Force Contract Nos. FA8702-15-D-0002, and FA870225DB003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

The opinions, findings, conclusions, and/or recommendations contained in this material are those of the author(s) and should not be construed as an official US Government position, policy, or decision, unless designated by other documentation.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution.

This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License (<https://creativecommons.org/licenses/by-nc/4.0/>). Requests for permission for non-licensed uses should be directed to the Software Engineering Institute at [permission@sei.cmu.edu](mailto:permission@sei.cmu.edu).

DM26-0329

## References

- Beck, K., Beedle, M., van Bennekum, A., Cockburn, A., Cunningham, W., Fowler, M., ... & Thomas, D. (2001). *Manifesto for agile software development*. <https://agilemanifesto.org/>
- Defense Innovation Board. (2019). *Software is never done: Refactoring the acquisition code for competitive advantage*. Department of Defense. [https://media.defense.gov/2019/apr/30/2002124828/-1/-1/0/softwareisneverdone\\_refactoringtheacquisitioncodeforcompetitiveadvantage\\_final.sw.ap.report.pdf](https://media.defense.gov/2019/apr/30/2002124828/-1/-1/0/softwareisneverdone_refactoringtheacquisitioncodeforcompetitiveadvantage_final.sw.ap.report.pdf)
- Department of Defense. (2020). *DoD Instruction 5000.87 Operation of the Software Acquisition Pathway*. <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500087p.PDF>
- Department of Defense. (2022). *DoD Instruction 5000.02 Operation of the Adaptive Acquisition Framework*. <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500002p.PDF>
- Department of War. (2025). *Acquisition transformation strategy: Rebuilding the arsenal of freedom*. <https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/ACQUISITION-TRANSFORMATION-STRATEGY.PDF>
- Fostering Reform and Government Efficiency in Defense (FoRGED) Act, S. 5618, 118th Cong. (2024). <https://www.congress.gov/bill/118th-congress/senate-bill/5618/text>



- Hegseth, P. (2025a). *Directing modern software acquisition to maximize lethality*.  
<https://media.defense.gov/2025/Mar/07/2003662943/-1/-1/1/DIRECTING-MODERN-SOFTWARE-ACQUISITION-TO-MAXIMIZE-LETHALITY.PDF>
- Hegseth, P. (2025b). *Transforming the defense acquisition system into the warfighting acquisition system to accelerate fielding of urgently needed capabilities to our warriors*.  
<https://media.defense.gov/2025/Nov/10/2003819439/-1/-1/1/TRANSFORMING-THE-DEFENSE-ACQUISITION-SYSTEM-INTO-THE-WARFIGHTING-ACQUISITION-SYSTEM-TO-ACCELERATE-FIELDING-OF-URGENTLY-NEEDED-CAPABILITIES-TO-OUR-WARRIORS.PDF>
- Hegseth, P. (2026). *Artificial intelligence strategy for the Department of War: Accelerating America's military AI dominance*. Department of War.  
<https://media.defense.gov/2026/Jan/12/2003855671/-1/-1/0/ARTIFICIAL-INTELLIGENCE-STRATEGY-FOR-THE-DEPARTMENT-OF-WAR.PDF>
- Hu, K. (2023, February 2). *ChatGPT sets record for fastest-growing user base - analyst note*. Reuters. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- Institute for Electrical and Electronics Engineers. (2025). *IEEE standard for the procurement of artificial intelligence and automated decision systems* (IEEE Std 3119-2025).  
<https://ieeexplore.ieee.org/document/11011522>
- Kennedy, B., Tyson, K., & Saks, E. (2023, February 15). *Public awareness of artificial intelligence in everyday activities: Limited enthusiasm in U.S. over AI's growing influence in daily life*. Pew Research Center.  
<https://www.pewresearch.org/science/2023/02/15/public-awareness-of-artificial-intelligence-in-everyday-activities/>
- Novak, W. E. (2021). *Artificial intelligence (AI) and machine learning (ML) acquisition and policy considerations* [White paper]. <https://www.sei.cmu.edu/library/artificial-intelligence-ai-and-machine-learning-ml-acquisition-and-policy-implications/>
- Rishel, P., Smith, C. J., O'Hearn, B., & Creel, R. C. (2025, August 5). Artificial intelligence in national security: Acquisition and integration. *Software Engineering Institute Blog*. Software Engineering Institute. <https://www.sei.cmu.edu/blog/artificial-intelligence-in-national-security-acquisition-and-integration/>
- Sidoti, O., & McClain, C. (2025, June 25). *34% of U.S. adults have used ChatGPT, about double the share in 2023*. Pew Research Center. <https://www.pewresearch.org/short-reads/2025/06/25/34-of-us-adults-have-used-chatgpt-about-double-the-share-in-2023/>
- Software Engineering Institute. (n.d.-a). *Pack your go bag*. Retrieved March 31, 2026, from SEI website: <https://www.sei.cmu.edu/projects/gobag/pack-your-go-bag/>
- Software Engineering Institute. (n.d.-b). *Software acquisition go bag*. Retrieved March 31, 2026, from SEI website: <https://www.sei.cmu.edu/projects/gobag/>
- Software Engineering Institute. (2025, December 12). *SWP essentials kit*.  
<https://www.sei.cmu.edu/library/swp-essentials-kit/>
- Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, J. S., Tambe, M., & Teller, A. (2016). Appendix I: A short history of AI. In *One*





*hundred year study on artificial intelligence: Report of the 2015–2016 Study Panel.*  
Stanford University. <https://ai100.stanford.edu/2016-report/appendix-i-short-history-ai>

Streamlining Procurement for Effective Execution and Delivery (SPEED) and National Defense Authorization Act for Fiscal Year (FY) 2026, H.R. 3838, 119th Cong. (2025).  
<https://www.congress.gov/bill/119th-congress/house-bill/3838>

Warfighting Acquisition University. (n.d.-a). *Capability development* [Glossary term]. WAU Glossary. Retrieved March 31, 2026, from <https://www.waru.edu/glossary/capability-development-document>

Warfighting Acquisition University. (n.d.-b). *Capability needs statement*. Software Acquisition. Retrieved March 31, 2026, from <https://aaf.dau.edu/aaf/software/cns/>

Warfighting Acquisition University. (n.d.-c). *Define capability needs*. Adaptive Acquisition Framework, Software Acquisition. Retrieved March 11, 2026, <https://aaf.dau.edu/aaf/software/define-capability-needs/>

Warfighting Acquisition University. (2020, October 2). *Planning phase*. Adaptive Acquisition Framework, Software Acquisition. <https://aaf.dau.edu/aaf/software/planning-phase/>

The White House. (2025, April 9). *Modernizing defense acquisitions and spurring innovation in the defense industrial base*. <https://www.whitehouse.gov/presidential-actions/2025/04/modernizing-defense-acquisitions-and-spurring-innovation-in-the-defense-industrial-base/>

Wicker, R. (2024, December 19). *Restoring freedom's forge: American innovation unleashed*. Defense Management Institute. <https://www.dmi-ida.org/knowledge-base-detail/Restoring-Freedoms-Forge-American-Innovation-Unleashed>



THIS PAGE INTENTIONALLY LEFT BLANK



## PANEL 3. ADAPTIVE LETHALITY: SYNCHRONIZING DIGITAL TWINS AND EMBEDDED ENGINEERING FOR RAPID FIELDING

Wednesday, May 6, 2026, 1145 – 1300 ET (0845 - 1000 PT)

### Panel Summary:

Modern conflict demands a departure from linear acquisition to meet the speed of evolving aerial and cyber-physical threats. This panel showcases three breakthrough methodologies for accelerating warfighting capabilities: the use of Cyber Digital Twins to self-heal critical infrastructure, embedded capability development to bridge the "valley of death" between prototypes and programs of record, and the rapid three-year integration of Airborne Directed Energy Weapons. Together, these researchers demonstrate how a "security-by-design" and "build-alongside" approach transforms experimental concepts into operational, cost-effective, and resilient weapon systems.

**Chair: Lieutenant General Frank Lozano, USA, Portfolio Acquisition Executive Fires (PAE Fires)**

### Presenters:

**Rapid Deployment of Airborne Directed Energy Weapons: A Strategic Solution to Evolving Aerial Threats – Ariel Dvorjetski, PEO for Weapons, Israeli Air Force**

**Embedded Capability Development: A Case Study in Rapid Missile Prototyping and Transition – Maj Dillon Pierce, PhD, Naval Postgraduate School / HQMC CD&I**

**Cyber Digital Twin-Informed Zero Trust: A Synergistic Framework for Securing Operational Technology in Defense Logistics Infrastructure – Dr. Barry A. Humphrey, Defense Logistics Agency**



**Lieutenant General Frank Lozano, USA**—is the Portfolio Acquisition Executive (PAE Fires), Redstone Arsenal, AL. He is the senior leader responsible for prioritizing and synchronizing acquisition functions across the Fires enterprise to deliver cutting edge warfighting capabilities to our Soldiers and joint warfighters. LTG Lozano assumed his current position in November 2025.

LTG Lozano graduated from Texas A&M University in 1993 and commissioned as a Second Lieutenant in the Armor Branch. He served as a Platoon Leader, Executive Officer, Cavalry Squadron Staff Officer, and Cavalry Troop Commander before transitioning into the Army Acquisition Corps in 2001. He graduated with an MBA from the University of Texas at Arlington and served with Lockheed Martin Missiles and Fire Control in Grand Prairie, TX in 2002 as part of the Training with Industry Program.

After completion of Command and General Staff College, LTG Lozano was assigned as the Assistant Product Manager for Project Manager Solider Weapons, Program Executive Officer (PEO) Solider, followed by an assignment as an Ammunition and Demolition Systems Acquisition Manager for Special Operations Command and the Army Research Development and Engineering Command.

In 2008, LTG Lozano was assigned as a Department of the Army System Coordinator for Tactical Missile Systems and Ballistic Missile Defense Systems. He was subsequently selected to be the Special Assistant to the Army's Vice Chief of Staff. In his position as the Special Assistant, he provided insight, advice, and counsel on Army acquisition programs crossing many different functional capability areas.

LTG Lozano commanded the Product Management Officer for Solider Protective Equipment, PEO Solider from 2011 until 2014. From 2014 to 2016 he was assigned to the Joint Staff, J-8 Capabilities and Acquisition Division. Upon graduation from the U.S. Army War College, he became the Project Manager for the Lower Tier Project Office, PEO Missiles and Space from 2017 until 2020, followed by an assignment as the Integrated Fires and Rapid Capability Office Program Manager. He was then assigned



as the ASA(ALT) Chief of Staff from 2021 to 2022, and from 2022 to 2025 he was appointed as the Program Executive Officer for PEO Missiles and Space.

LTG Lozano's operational and combat experience include deployments to Bosnia, Kuwait, and Iraq. His awards and decorations include the Parachutist Badge, Ranger Tab, Legion of Merit, Bronze Star Medal, joint Service Commendation Medal, the NATO Service Medal, the Army Staff Identification Badge, and the Joint Staff Identification Badge. He is certified in Program Management, Contractive, System Research, Planning and Engineering, and System Test career fields.



# **Rapid Deployment of Airborne Directed Energy Weapons: A Strategic Solution to Evolving Aerial Threats**

**Ariel Dvorjetski**—PEO for Weapon, Israel Air Force [arieldvorjetski@idf.il]

**Raymond D. Jones**—Professor of the Practice, Department of Acquisition, Finance and Manpower, Naval Postgraduate School. [rdjone1@nps.edu]

**Bonnie Johnson**—Researcher, Naval Postgraduate School. [bwjohnso@nps.edu]

## **Abstract**

Modern aerial warfare is increasingly characterized by low-cost, high-volume threats such as unmanned aerial vehicles (UAVs) and cruise missiles, which challenge the economic and operational sustainability of traditional kinetic air defense systems. This paper examines the feasibility of rapidly deploying airborne Directed Energy Weapons (DEWs), specifically high-TRL solid-state lasers, as a strategic response to this imbalance. The research integrates technical analysis, operational scenario modeling, and acquisition strategy to define a realistic pathway for deployment within a three-year timeframe. A 30–50 kW class laser system is identified as the optimal balance between operational effectiveness and Size, Weight, Power, and Cooling (SWaP-C) constraints, enabling multiple engagements per sortie while remaining compatible with existing airborne platforms. Scenario-based analysis demonstrates that such systems can effectively counter high-volume UAV threats and enhance layered defense architectures by improving cost-exchange ratios and operational persistence. The study further argues that traditional acquisition processes are incompatible with current threat timelines and proposes the Middle Tier of Acquisition (MTA) as a critical enabler for rapid fielding. The findings indicate that airborne DEWs are not only technically feasible but strategically necessary, providing a scalable and sustainable solution for modern air defense.

## **Introduction**

Modern aerial warfare is undergoing a structural transformation driven not by technological superiority alone, but by the growing dominance of low-cost, high-volume threats. Unmanned aerial vehicles (UAVs) and cruise missiles are increasingly employed as scalable, accessible, and economically efficient tools of warfare. As highlighted by Ressler (2024), these systems enable adversaries to project power at relatively low cost while exploiting vulnerabilities in traditional defense architectures. This shift has exposed a fundamental weakness in current air defense strategies: the inability to sustain a favorable cost-exchange ratio over time.

Conventional air defense systems rely heavily on kinetic interceptors such as surface-to-air missiles. While these systems are highly effective in terms of precision and lethality, they are constrained by both cost and inventory. Each engagement carries a significant financial burden, often reaching hundreds of thousands or even millions of dollars per interceptor. When employed against low-cost threats, this creates an economic imbalance that is strategically unsustainable. As demonstrated in recent operational contexts, defenders are increasingly forced into a reactive posture in which success in interception does not translate into long-term strategic advantage (Sayler et al., 2024).

This challenge is further exacerbated by the emergence of swarm tactics, where large numbers of inexpensive UAVs are used to saturate and overwhelm defense systems. Traditional interceptors, limited by magazine depth and reload constraints, are not designed to counter such high-volume attacks efficiently. As a result, the current defense paradigm is being stretched beyond its intended limits, requiring new approaches that can scale both operationally and economically.








**Figure 1. Iran Deadly Drone Arsenal Reach**  
(Middle East Council on Global Affairs, 2024)

Directed Energy Weapons (DEWs), and specifically solid-state laser systems, offer a promising alternative to this paradigm. Unlike kinetic interceptors, DEWs deliver energy directly to a target at the speed of light, enabling near-instantaneous engagement without the need for physical munitions. This fundamentally changes the cost structure of air defense, as the marginal cost per engagement is reduced to the cost of energy consumption, which is negligible compared to traditional interceptors (Johnson et al., 2023). Additionally, DEWs provide the potential for a “deep magazine,” limited primarily by onboard power generation rather than physical inventory.

The growing maturity of solid-state laser technology further strengthens the case for their operational deployment. Advances in beam control, power generation, and system miniaturization have brought these systems to a Technology Readiness Level (TRL) that supports near-term integration into military platforms (Blau et al., 2018; Weinberg, 2023). Real-world demonstrations, including systems such as ODIN and HELIOS, have already shown the feasibility of using directed energy to counter aerial threats in operational environments (Johnson, 2024; Saylor et al., 2024).

Despite these advancements, the transition from experimental systems to operational capability remains a significant challenge. Traditional defense acquisition processes are often characterized by long development cycles, complex requirements, and extensive oversight, which can delay the fielding of new technologies. In the context of rapidly evolving threats, such timelines are no longer acceptable. The central premise of this research is therefore that airborne DEWs must be deployed within a compressed timeframe—specifically, within three years—to remain operationally relevant.

To achieve this objective, the research adopts a pragmatic approach that prioritizes mature technologies and focuses on integration feasibility. Among the various DEW technologies evaluated, solid-state lasers emerge as the most suitable candidate due to their relatively high TRL, scalability, and compatibility with airborne platforms. The study identifies a 30–50 kW class laser system as a realistic “working point,” balancing operational effectiveness with the constraints imposed by size, weight, power, and cooling (SWaP-C) requirements. This configuration is expected to enable multiple engagements per sortie, providing meaningful operational impact while remaining feasible within existing platform limitations.

A key conceptual tension explored in this research is the balance between “technology push” and “requirement pull.” While cutting-edge technologies offer significant potential, their development timelines often exceed operational needs. Conversely, strictly requirement-driven approaches may limit innovation. This study proposes a hybrid approach, leveraging mature technologies while adapting them rapidly to meet urgent operational requirements. This perspective aligns with broader strategic frameworks that emphasize adaptability and iterative development in complex environments (Reeves et al., 2015).

In addition to technological considerations, the research addresses the broader system-level challenges associated with deploying airborne DEWs. These include integration into existing command-and-control architectures, development of operational doctrines, and implementation of safety and regulatory frameworks (Cannin, 2021). The complexity of these challenges highlights that the deployment of DEWs is not solely a technical problem, but also an organizational and strategic one.

Ultimately, this research seeks to answer a central question: How can DEWs be effectively integrated onto airborne platforms within a three-year timeframe to counter modern aerial threats in a cost-effective and sustainable manner? By combining technical analysis, operational scenario modeling, and acquisition strategy considerations, the study aims to provide a realistic and actionable framework for the rapid deployment of airborne DEWs.

The significance of this work extends beyond the specific application of directed energy. It represents a broader shift in how defense systems are conceptualized and developed in an era defined by asymmetric threats and rapid technological change. Airborne DEWs are not merely an incremental improvement—they are a potential enabler of a new defense paradigm, one that prioritizes sustainability, scalability, and strategic adaptability.

## **Problem Context**

The contemporary defense environment is increasingly shaped by the proliferation of low-cost, high-volume aerial threats, particularly unmanned aerial vehicles (UAVs) and cruise missiles. These systems are not inherently advanced, but their effectiveness lies in their scalability, accessibility, and ability to exploit structural weaknesses in existing defense architectures. As highlighted by Rassler (2024), the evolution of such threats reflects a broader shift toward cost-efficient and distributed methods of warfare, where quantity and persistence can outweigh technological superiority.



### Shahed-136 Drone ("Geran-2")

An Iranian-made suicide drone imported by Russia for use in its invasion of Ukraine. Rebranded as "Geran-2," or Geranium-2.




**Figure 2. Iran Shahed 136 Key Performance Parameters**

Traditional air defense systems were not designed for this type of threat environment. Their architecture is optimized for intercepting high-value targets using precision-guided kinetic interceptors. While effective against limited and sophisticated threats, these systems face significant limitations when confronted with large-scale, low-cost attacks. The fundamental issue is not the performance of individual interceptors, but the **economic and operational model** underlying their use.

Kinetic interceptors are expensive, finite, and logistically demanding. Each engagement consumes a physical asset that must be replenished, transported, and maintained. In contrast, UAVs and similar threats can be produced rapidly, at low cost, and deployed in large numbers. This creates a structural asymmetry in which the defender is forced to expend disproportionate resources to counter relatively inexpensive threats. As noted in Department of Defense assessments, this imbalance becomes particularly critical in scenarios involving swarm tactics, where saturation effects can overwhelm even advanced defense systems (Saylor et al., 2024).

The operational implications of this imbalance are significant. First, there is a limitation in **magazine depth**. Aircraft and ground-based systems can carry only a finite number of interceptors, which constrains their ability to sustain prolonged engagements. Second, there is a constraint on **engagement rate**, as kinetic systems require time for targeting, launch, and potential re-engagement. Third, and most critically, there is a **cost-exchange problem**, where the cumulative cost of defense exceeds the cost imposed on the attacker.

This dynamic is already evident in recent operational contexts. The increasing use of UAV swarms and missile barrages has demonstrated the vulnerability of traditional systems to saturation attacks, particularly when defenders are required to allocate high-value interceptors against low-value targets. As Johnson (2024) observes in the context of recent maritime operations, directed energy systems have begun to emerge as a response to precisely this challenge, offering a more sustainable engagement model.

Directed Energy Weapons (DEWs), particularly solid-state lasers, directly address the core limitations of kinetic systems. Unlike traditional interceptors, DEWs do not rely on physical

munitions and therefore are not constrained by magazine depth in the conventional sense. Instead, their operational limits are defined by power generation and thermal management. This enables a fundamentally different engagement paradigm, where multiple targets can be engaged sequentially with minimal marginal cost (Johnson et al., 2023).

However, while DEWs offer clear theoretical advantages, their practical implementation introduces a new set of challenges. The airborne environment imposes strict constraints on size, weight, power, and cooling (SWaP-C), which directly impact the feasibility of integrating laser systems onto existing platforms. As demonstrated in prior research on laser integration, achieving a balance between system performance and platform limitations is one of the most critical design challenges (Blau et al., 2018).

In addition to technical constraints, there are significant operational and doctrinal considerations. DEWs must be integrated into existing command-and-control architectures and used in conjunction with kinetic systems to form a layered defense. This requires the development of new engagement concepts that leverage the strengths of both approaches. Furthermore, safety considerations—such as preventing unintended exposure to friendly forces or civilian assets—must be addressed through robust control mechanisms and operational procedures (Cannin, 2021).

Another critical dimension of the problem is the timeline. Traditional defense acquisition processes are not aligned with the pace of threat evolution. Programs often require extensive development cycles, which delay the transition from prototype to operational capability. As emphasized in acquisition and innovation literature, bridging this gap requires more adaptive and flexible approaches that can accelerate development while managing risk (GAO, 2023; Reeves et al., 2015).

This research frames the problem as a **multi-dimensional challenge** involving:

- Economic imbalance (cost-exchange ratio)
- Operational constraints (magazine depth, engagement rate)
- Technical limitations (SWaP-C integration)
- Organizational barriers (acquisition timelines and processes)

The central argument is that addressing only one of these dimensions is insufficient. A viable solution must operate across all four simultaneously.

Therefore, the focus of this study is not simply on evaluating DEWs as a technology, but on assessing their feasibility as an **integrated operational and acquisition solution**. Specifically, the research examines whether a high-TRL solid-state laser system can be deployed on an airborne platform within a three-year timeframe while maintaining operational effectiveness and economic viability.

By situating the problem within this broader systemic context, the research moves beyond a purely technical discussion and establishes the foundation for a comprehensive solution. The following sections build on this foundation by examining the technological options, operational scenarios, and acquisition pathways that can enable the rapid deployment of airborne DEWs in response to evolving aerial threats.

## Technological Comparison and Selection

The selection of an appropriate Directed Energy Weapon (DEW) technology for airborne deployment is not merely a technical decision, but a strategic one. It directly determines the feasibility of meeting the central requirement of this research: achieving an operational



capability within a three-year timeframe. While multiple DEW technologies have been developed over the past decades, their suitability varies significantly when evaluated against criteria such as technological maturity, integration feasibility, operational effectiveness, and logistical sustainability.

Three primary categories of laser-based DEWs are considered in this analysis: **chemical lasers**, **free electron lasers (FELs)**, and **solid-state lasers (SSLs)**. Each represents a different point along the spectrum of power, complexity, and maturity.

Chemical lasers were among the earliest high-energy laser systems developed for military applications. They are capable of generating very high power levels and have demonstrated effectiveness in controlled environments. However, their operational use is constrained by significant logistical and safety challenges. Chemical lasers rely on hazardous reactants, require complex handling procedures, and impose substantial weight and volume burdens on the platform. These characteristics make them unsuitable for airborne integration, particularly under the strict Size, Weight, Power, and Cooling (SWaP-C) constraints identified in this research. As a result, despite their high output potential, chemical lasers are not considered viable for rapid deployment scenarios (Blau et al., 2018).

FELs represent a more advanced and flexible technological approach. They offer the potential for extremely high power output and tunable wavelengths, enabling adaptability across different operational conditions. However, FELs remain at a relatively low Technology Readiness Level (TRL) and require large, complex infrastructures that are incompatible with airborne platforms. Their size, power requirements, and developmental maturity place them firmly outside the scope of a three-year deployment timeline. While FELs may play a role in future high-power applications, they do not meet the near-term operational requirements defined in this study (Weinberg, 2023).

In contrast, SSLs provide a balanced solution that aligns closely with both technical and operational constraints. SSLs have reached a relatively high level of maturity, with many systems demonstrating TRL levels between 6 and 8. This maturity is reflected in recent operational demonstrations and ongoing programs, where SSLs have been successfully integrated into land-based and maritime platforms (Johnson, 2024; Saylor et al., 2024). Their architecture allows for modular scaling of power, improved reliability, and reduced logistical complexity compared to chemical systems.

From an operational perspective, SSLs offer several critical advantages. First, they enable **low-cost engagements**, as the primary resource consumed is electrical energy rather than physical munitions. Second, they provide **speed-of-light engagement**, which eliminates time-of-flight constraints and improves effectiveness against fast or maneuvering targets. Third, they support **scalability**, allowing system performance to be adjusted through incremental increases in power output (Johnson et al., 2023).

However, the selection of SSLs does not eliminate all challenges. Their integration into airborne platforms introduces significant constraints related to power generation, thermal management, and beam quality. Atmospheric effects, including turbulence and absorption, can degrade laser performance and must be mitigated through advanced beam control and adaptive optics (Sabatini et al., 2015; Shi et al., 2023). These limitations highlight the importance of selecting not only the appropriate technology type, but also the appropriate **power level**.

A key contribution of this research is the identification of a realistic “working point” for airborne DEWs. Rather than pursuing maximum power output, which would exceed platform constraints, the study focuses on achieving a balance between effectiveness and feasibility. The analysis indicates that a **30–50 kW class solid-state laser** represents the optimal range for





near-term deployment. Within this range, the system can deliver meaningful operational effects—such as neutralizing UAVs within tactically relevant engagement times—while remaining compatible with the SWaP-C limitations of existing aircraft.

This approach reflects a deliberate shift from a “maximum capability” mindset to a **“sufficient capability” framework**. The objective is not to replace kinetic systems across all threat types, but to provide a complementary capability that addresses the specific challenge of high-volume, low-cost threats. Even with limited dwell time and engagement range, a system in this power class can significantly reduce the burden on traditional interceptors by absorbing a substantial portion of incoming threats.

The selection of solid-state lasers, therefore, is not based solely on their theoretical advantages, but on their ability to meet a complex set of constraints simultaneously. They represent the only DEW technology that satisfies the combined requirements of:

- High technological maturity
- Feasible airborne integration
- Operational effectiveness against UAV-class threats
- Compatibility with a three-year deployment timeline

This conclusion reinforces the broader argument of the research: that rapid deployment is achievable only by aligning technological ambition with practical constraints. By focusing on solid-state lasers within a defined power range, the study establishes a credible foundation for the subsequent analysis of system integration, operational performance, and acquisition strategy.

## Key Performance Parameters (KPPs) and System Design

The transition from conceptual analysis to an operationally viable solution requires the definition of clear and achievable Key Performance Parameters (KPPs). These parameters serve as the bridge between technological capability and mission effectiveness, ensuring that the proposed Directed Energy Weapon (DEW) system is not only theoretically sound but also practically deployable within the constraints identified throughout this research.

Rather than optimizing for maximum performance, this study adopts a **constraint-driven design approach**, where system parameters are derived from the intersection of operational requirements and platform limitations. This approach reflects the central premise of the research: that rapid deployment within a three-year timeframe requires prioritizing feasibility, integration, and sufficiency over technological perfection.

The primary constraint shaping system design is the airborne environment, which imposes strict limitations on size, weight, power, and cooling (SWaP-C). As demonstrated in prior studies on laser integration, these constraints are often the decisive factor in determining whether a system can transition from prototype to operational capability (Blau et al., 2018). Consequently, the definition of KPPs must begin with an understanding of what can realistically be integrated onto existing platforms without extensive redesign.

Based on this framework, the research identifies a **30–50 kW class solid-state laser** as the optimal power range for airborne deployment. This range represents a critical balance point. At lower power levels, the system’s ability to achieve reliable target neutralization within operationally relevant timeframes is limited. At higher power levels, the demands on power generation, thermal management, and structural integration exceed what can be accommodated within current airborne platforms.



Within this range, the study converges on a **30 kW system as the primary working point**, providing a realistic and testable configuration. This selection is supported by both engineering considerations and operational modeling, which indicate that such a system can achieve meaningful effects against UAV-class threats while remaining compatible with SWaP-C constraints.

Several KPPs emerge from this analysis:

- First, **laser power output**, defined at approximately 30 kW, represents the minimum threshold required to achieve effective engagement against small UAVs within tactically acceptable dwell times. This level of power enables the system to deliver sufficient energy to critical components of the target, resulting in mission kill rather than complete destruction, which aligns with the objective of efficient threat neutralization.
- Second, **effective engagement range**, estimated at approximately 5–7 kilometers, reflects the balance between beam propagation limitations and operational relevance. Atmospheric effects such as absorption, scattering, and turbulence impose constraints on laser performance, requiring careful optimization of engagement distances (Sabatini et al., 2015). Within this range, the system can operate effectively while maintaining integration feasibility.
- Third, **engagement capacity per sortie**, estimated at approximately 10–20 targets, represents a critical metric for evaluating operational value. Unlike kinetic systems, where each engagement consumes a physical interceptor, DEWs allow for multiple sequential engagements limited primarily by power availability and thermal constraints. This capability significantly enhances mission endurance and reduces dependence on resupply.
- Fourth, **system weight**, constrained to approximately 3,000 pounds or less, is essential for maintaining aircraft performance and avoiding major structural modifications. This requirement reinforces the need for a pod-based solution, where the DEW system is integrated as a modular unit rather than as a fully embedded subsystem.
- Fifth, **power and thermal management requirements** define the supporting infrastructure needed to sustain laser operation. These include onboard power generation, energy storage, and cooling systems capable of maintaining performance under continuous operation. Advances in energy storage and thermal management technologies are critical enablers in this context (Kim, 2024; Sylvester, 2014).

Together, these parameters define a system that is not optimized for maximum destructive capability, but for **operational relevance within real-world constraints**. This distinction is essential. The objective is not to replace kinetic interceptors across all threat categories, but to provide a complementary capability that addresses the specific challenge of high-volume, low-cost threats.






**Figure 3. Total Weight of Airborne Laser Pod: Weight Assessment vs. Operational Range for Different Power and Number of Enemy Threats in One Sortie**

To validate the operational implications of these KPPs, the research incorporates scenario-based modeling that simulates engagements between an airborne DEW-equipped platform and multiple UAV threats. The results indicate that a system within the defined parameter space can successfully neutralize multiple targets within a single sortie, with engagement times that are consistent with real-world operational requirements. These findings support the conclusion that even a relatively moderate-power system can deliver disproportionate operational value when evaluated within the broader context of cost efficiency and system sustainability.

An important insight emerging from this analysis is the concept of **“sufficient lethality.”** Unlike kinetic systems, which are often designed for complete target destruction, DEWs can achieve mission success through partial damage to critical components. This reduces the energy required per engagement and allows the system to maximize the number of targets addressed within a given mission.

Furthermore, the defined KPPs align closely with the constraints of rapid acquisition and deployment. By limiting system complexity and focusing on achievable performance thresholds, the design supports iterative development and testing, which are essential for meeting the three-year timeline. This approach is consistent with adaptive acquisition strategies that emphasize incremental capability delivery over long-term optimization (Reeves et al., 2015).

In summary, the KPPs defined in this research represent a carefully balanced solution that integrates technological capability, operational requirements, and platform constraints. The identification of a 30 kW-class solid-state laser as the system’s core element provides a credible and actionable foundation for the subsequent analysis of integration, operational employment, and acquisition strategy.

## Operational Scenarios and System Validation

The definition of Key Performance Parameters (KPPs) provides a necessary foundation for system design, but it is insufficient on its own to establish operational relevance. To evaluate the effectiveness and scalability of airborne Directed Energy Weapons (DEWs), this research employs a scenario-based approach, examining how the proposed system performs under realistic and diverse operational conditions. This approach aligns with established principles of defense analysis, which emphasize the importance of contextual validation in complex and uncertain environments (Danzig, 2011).

Two distinct operational scenarios are used to test the robustness of the proposed solution. These scenarios are deliberately selected to represent different ends of the conflict spectrum: a localized, high-intensity environment characterized by UAV swarm threats and a broader, strategic environment involving long-range missile and UAV campaigns. Together, they provide a comprehensive framework for assessing both tactical effectiveness and strategic scalability.

The first scenario reflects a short-range, high-volume threat environment, similar to those observed in conflicts involving non-state actors such as Hezbollah. In this context, UAV swarms and short-range missiles are used to saturate defense systems through continuous and distributed attacks. These threats are characterized by low individual cost, high frequency, and the ability to exploit gaps in air defense coverage. Traditional kinetic systems struggle in such environments due to their limited magazine depth and high cost per engagement.

Within this scenario, airborne DEWs provide a critical advantage. By enabling multiple engagements per sortie with minimal marginal cost, the system can absorb a significant portion of the incoming threat volume. The ability to engage targets sequentially, without reliance on physical munitions, allows the platform to maintain operational effectiveness even during prolonged engagements. This directly addresses one of the core limitations of kinetic systems: the inability to sustain high engagement rates against large numbers of targets.

The second scenario represents a long-range, state-level conflict, such as a large-scale missile and UAV campaign driven by actors like Iran, with potential relevance to future operations in theaters such as INDOPACOM. In this environment, threats include not only UAVs but also more advanced systems, including long-range cruise missiles and coordinated strike packages. The scale and complexity of such scenarios require a defense architecture that is both layered and scalable.

In this context, airborne DEWs are not expected to replace kinetic systems, but to enhance them. Their role is to act as a **force multiplier**, reducing the burden on high-value interceptors by neutralizing lower-tier threats. By filtering out high-volume, low-cost targets, DEWs allow kinetic systems to focus on more complex and high-priority threats. This layered approach increases overall system efficiency and resilience, particularly in scenarios where resource allocation becomes a critical factor.

To support these scenario analyses, the research incorporates operational modeling of engagement dynamics. The modeling focuses on interactions between an airborne DEW-equipped platform and multiple UAV threats within a defined battlespace. Parameters such as target speed, engagement range, and dwell time are used to simulate realistic conditions. The results indicate that a 30 kW-class system can successfully neutralize multiple UAVs within a single engagement window, with dwell times that are consistent with known laser performance characteristics (Shi et al., 2023).

A key finding from this modeling is the system's ability to maintain a favorable engagement tempo. Unlike kinetic systems, which are constrained by reload cycles and



interceptor availability, the DEW system can transition rapidly between targets. This enables it to handle sequential engagements efficiently, making it particularly effective against swarm-based threats. The cumulative effect is a significant increase in the number of targets that can be addressed per sortie, reinforcing the system's role as a scalable solution.




**Figure 4. Dynamic Simulation of Jetfighter Against 3 UAV Threat in a Square Area of 10 km<sup>2</sup>**

Another important aspect of the analysis is **fleet-level scalability**. While a single aircraft equipped with a DEW system can provide localized defense, the full operational value emerges when multiple platforms are integrated into a coordinated network. In such a configuration, DEW-equipped aircraft can be distributed across the battlespace, providing overlapping coverage and enhancing overall system redundancy. This distributed approach aligns with modern concepts of multi-domain operations, where flexibility and adaptability are critical to mission success.

The research also examines the implications of using a single DEW configuration across multiple scenarios. Given the constraints of the three-year deployment timeline, developing separate systems for different threat environments is not feasible. Instead, the study emphasizes the importance of a **multi-role system** capable of adapting to both localized and strategic contexts. The 30–50 kW class system identified in this research meets this requirement, offering sufficient flexibility to operate effectively across a range of scenarios.

In addition to operational effectiveness, the scenarios highlight the importance of integration with command-and-control (C2) systems. Effective use of DEWs requires real-time target identification, prioritization, and engagement coordination. This necessitates seamless integration with existing sensor networks and decision-making frameworks. Open architecture interfaces and standardized communication protocols are essential to enabling this integration and ensuring that DEWs can function as part of a cohesive defense system (Lazov et al., 2021).

Finally, the scenario analysis reinforces the broader strategic value of airborne DEWs. Their ability to provide persistent, low-cost engagement capability enhances not only tactical effectiveness but also overall system resilience. By reducing dependence on finite resources



and enabling sustained operations, DEWs contribute to a more sustainable defense posture in the face of evolving threats.

In summary, the operational scenarios demonstrate that the proposed airborne DEW system is not only technically feasible but also operationally relevant across a range of conflict environments. The combination of tactical effectiveness, scalability, and integration potential supports the conclusion that DEWs can play a central role in future air defense architectures, provided they are deployed within an appropriate strategic and organizational framework.

## Acquisition Strategy and Rapid Deployment Framework

The feasibility of airborne Directed Energy Weapons (DEWs) is not determined solely by technological maturity or operational effectiveness. The decisive factor is whether these capabilities can be **acquired, integrated, and fielded within an operationally relevant timeframe**. In the context of rapidly evolving aerial threats, traditional defense acquisition processes are no longer sufficient. The central challenge, therefore, is not only technological—but organizational and procedural.

Conventional acquisition frameworks within the Department of Defense are designed to minimize risk through extensive requirements definition, sequential development phases, and rigorous oversight. While this approach has proven effective for large, complex systems, it typically results in development timelines exceeding a decade. Such timelines are fundamentally incompatible with the pace at which threats such as UAV swarms and low-cost missile systems are evolving. As highlighted by the Government Accountability Office (GAO, 2023), one of the most persistent challenges in defense innovation is the “valley of death,” where promising technologies fail to transition into operational capability due to bureaucratic and programmatic barriers.

To address this gap, this research proposes the use of the **Middle Tier of Acquisition (MTA)** as the primary acquisition pathway for airborne DEWs. The MTA framework is specifically designed to accelerate the development and fielding of capabilities by reducing procedural constraints and enabling iterative development. It provides two complementary pathways: rapid prototyping and rapid fielding. Together, these pathways offer a structured yet flexible approach to achieving operational capability within a compressed timeline.

The rapid prototyping phase focuses on developing and demonstrating a functional system within a 12–24-month period. In the context of airborne DEWs, this phase includes the design and construction of a pod-based laser system, integration with a host platform, and initial ground-based testing. The emphasis is on validating core functionalities—power generation, beam control, and thermal management—rather than achieving full operational optimization. This aligns with the broader principle of prioritizing **minimum viable capability** over maximum performance.

The rapid fielding phase builds on the results of prototyping and focuses on transitioning the system into operational use. This includes flight testing, refinement of system integration, development of operational procedures, and initial deployment within an operational unit. The objective is to achieve Initial Operational Capability (IOC) within a total timeframe of approximately three years. This accelerated timeline is consistent with adaptive acquisition strategies that emphasize speed, iteration, and user feedback (Reeves et al., 2015).

A key enabler of this approach is the adoption of a **modular, pod-based architecture**. By integrating the DEW system as an external pod rather than a fully embedded subsystem, the program avoids the need for extensive aircraft redesign. This significantly reduces integration risk, shortens development timelines, and allows for parallel testing across multiple platforms.



The modular approach also supports scalability, enabling future upgrades without requiring major structural modifications.




**Figure 5. A Pragmatic Approach to De-Risking Complex Integration**

Another critical element of the acquisition strategy is the use of **iterative development and spiral upgrades**. Rather than attempting to deliver a fully optimized system in a single step, the program evolves through successive increments. Each iteration introduces improvements in power, efficiency, and integration while maintaining operational capability throughout the process. This approach reduces risk by allowing issues to be identified and addressed early, and it aligns with best practices in modern defense innovation.

Risk management is a central component of the acquisition framework. The primary risks associated with airborne DEWs include power generation, thermal management, beam control, and system integration. These risks are mitigated through a combination of technology selection, system design, and testing strategy. For example, the use of high-TRL solid-state lasers reduces technological uncertainty, while the pod-based architecture minimizes integration complexity. Advanced modeling tools, including digital twins and predictive analytics, further support risk reduction by enabling early identification of potential issues and optimization of system performance (Weinberg, 2023).

In addition to technical risks, the acquisition strategy must address organizational and regulatory challenges. Effective coordination between military stakeholders, industry partners, and regulatory bodies is essential to ensure alignment of requirements, resources, and timelines. Open architecture principles play a key role in this context, enabling interoperability and facilitating integration with existing command-and-control systems (Lazov et al., 2021).

The integration of **Systemic Operational Design (SOD)** principles provides an additional layer of strategic coherence to the acquisition process. SOD emphasizes the importance of understanding complex systems as dynamic and interconnected, where tensions and constraints can be leveraged to drive innovation. In this research, the three-year deployment timeline is framed as a “productive tension,” forcing prioritization, accelerating decision-making, and encouraging creative problem-solving (Graicer, 2017). This perspective shifts the focus from eliminating constraints to using them as drivers of efficiency and innovation.

An important consideration within the acquisition strategy is whether to adopt an intermediate step, such as integrating the DEW system onto a larger transport aircraft before transitioning to fighter platforms. While this approach may reduce initial risk by providing a more permissive integration environment, it also introduces potential delays and additional program

complexity. This research argues that, given the urgency of the operational requirement, a direct path to fighter integration—supported by modular design and incremental testing—offers a more effective balance between risk and speed.

Ultimately, the proposed acquisition framework reflects a fundamental shift in how defense capabilities are developed and fielded. It moves away from a linear, risk-averse model toward a more agile and adaptive approach that prioritizes speed, flexibility, and operational relevance. This shift is not optional; it is a necessary response to a threat environment that evolves faster than traditional acquisition processes can accommodate.

In summary, the rapid deployment of airborne DEWs is achievable not only because the technology is sufficiently mature, but because acquisition frameworks such as MTA provide the means to translate that maturity into operational capability. By aligning technology, system design, and acquisition strategy within a coherent framework, this research demonstrates that a three-year deployment timeline is both realistic and attainable.




Figure 6. Three-Year Path to Capability

## Strategic Implications and Conclusions

The rapid deployment of airborne Directed Energy Weapons (DEWs) represents more than a technological advancement; it reflects a fundamental shift in the logic of air defense. The findings of this research demonstrate that DEWs are not merely an incremental improvement to existing systems, but a capability that alters the underlying balance between offense and defense in modern aerial warfare.

At the core of this transformation is the **cost-exchange dynamic**. Traditional kinetic systems impose a structural disadvantage on defenders, forcing them to expend high-value interceptors against low-cost threats. Over time, this creates an unsustainable model, where even successful defenses erode strategic resources. As highlighted in prior analyses, this imbalance is already evident in operational environments characterized by UAV swarms and missile saturation attacks (Sayler et al., 2024).

Airborne DEWs directly address this challenge by fundamentally reducing the marginal cost of engagement. By shifting the constraint from inventory to energy, they enable a more sustainable defense posture, where the defender can match or exceed the attacker's scale without proportional increases in cost. This transition is not simply economic—it is strategic. It

restores the defender's ability to absorb high-volume attacks without degrading long-term readiness.

Beyond cost considerations, DEWs introduce a new level of **operational persistence**. Unlike kinetic systems, which are limited by magazine depth and reload cycles, DEWs can sustain engagements as long as sufficient power is available. This creates a continuous defensive capability that is particularly effective against swarm-based threats. In this sense, DEWs transform air defense from a sequence of discrete engagements into a **persistent engagement environment**, where threats can be addressed dynamically and continuously.

Another important implication is the impact on **force design and layered defense architectures**. The integration of DEWs enables a more efficient allocation of resources across different threat tiers. High-cost kinetic interceptors can be reserved for complex, high-value targets, while DEWs handle high-volume, low-cost threats. This layered approach increases overall system efficiency and enhances resilience against saturation attacks. It also aligns with broader concepts of multi-domain operations, where flexibility and adaptability are critical to mission success.

From a technological perspective, the research confirms that the deployment of airborne DEWs does not require waiting for breakthrough innovations. High-TRL solid-state lasers already provide sufficient capability to deliver meaningful operational effects. The identification of a 30–50 kW class system as a viable “working point” demonstrates that **sufficient capability**, rather than maximum capability, is the key to rapid deployment. This insight challenges traditional development models that prioritize long-term optimization over near-term applicability.

Equally important is the role of **acquisition strategy** in enabling this transformation. The research shows that traditional acquisition processes are incompatible with the urgency of the current threat environment. The adoption of frameworks such as the Middle Tier of Acquisition (MTA) is not simply an efficiency improvement—it is a prerequisite for operational relevance. By enabling rapid prototyping, iterative development, and early fielding, MTA allows the defense community to align technological maturity with operational timelines (GAO, 2023).

The integration of Systemic Operational Design (SOD) principles further reinforces this approach by framing constraints—such as the three-year deployment timeline—not as limitations, but as drivers of innovation. This perspective encourages a shift from risk avoidance to **risk management through iteration**, enabling faster and more adaptive development processes (Graicer, 2017).








**Figure 7: Integrated Strategy for Rapid Deployment and Cost-Effective Defense of DEWs**

At the strategic level, the deployment of airborne DEWs signals a broader transition toward **sustainable defense systems**. In an era defined by asymmetric threats and rapid technological change, sustainability becomes a critical factor in maintaining long-term superiority. DEWs contribute to this by reducing logistical burdens, increasing operational endurance, and enabling scalable responses to evolving threats.

However, the successful implementation of airborne DEWs also requires addressing several ongoing challenges. Technical issues related to power generation, thermal management, and beam control must continue to be refined. Operational doctrines must evolve to incorporate DEWs into existing frameworks, ensuring effective coordination with kinetic systems. Regulatory and safety considerations must be carefully managed to prevent unintended consequences, particularly in complex operational environments (Cannin, 2021).

Despite these challenges, the overall conclusion of this research is clear. The rapid deployment of airborne DEWs is both **feasible and strategically necessary**. The convergence of mature technology, evolving threats, and adaptive acquisition frameworks creates a unique opportunity to redefine air defense within a relatively short timeframe.

The implications extend beyond the specific application of DEWs. They reflect a broader lesson for defense innovation: The ability to align technology, operations, and acquisition processes is more critical than any single technological breakthrough. In this context, airborne DEWs serve as a case study in how such alignment can be achieved.

In conclusion, this research demonstrates that a three-year pathway to operational airborne DEWs is realistic, provided that the focus remains on mature technologies, disciplined system design, and accelerated acquisition strategies. More importantly, it shows that the adoption of DEWs has the potential to reshape the future of aerial defense, shifting it toward a model that is more resilient, scalable, and economically sustainable.

The challenge ahead is not whether this transformation is possible—but whether organizations can move fast enough to realize it.



## References

- Air Force Research Laboratory. (2020). *Self-protect high energy laser demonstrator (SHiELD)* (AFMC-2020-0021). <https://www.afresearchlab.com>
- Blau, J., Johnson, B., Green, J., Michnewich, D., Gildemeyer, S., Hager, D., Liensdorf, D., Malone, A., & Mugerditchian, K. (2018). *Shipboard laser analysis: SWAP-C and operational capabilities*. Naval Postgraduate School. <https://hdl.handle.net/10945/64355>
- Cannin, A. J. (2021). *Directed energy weapons: Ethical implementation obstacles* [Master's thesis, Naval Postgraduate School]. <https://hdl.handle.net/10945/68701>
- Danzig, R. (2011). *Driving in the dark: Ten propositions about prediction and national security*. Center for a New American Security. <https://www.jstor.org>
- GAO. (2023). *Directed energy weapons: DOD should focus on transition planning* (GAO-23-105868). <https://www.gao.gov/products/gao-23-105868>
- Graicer, O. (2017). Self-disruption: Seizing the high ground of systemic operational design (SOD). *Journal of Military and Strategic Studies*, 17(4), 30–37. <https://jmss.org/article/view/58253/pdf>
- Johnson, B. (2024). *The coming of age of directed energy weapons and the Red Sea crisis*. Center for International Maritime Security. <https://cimsec.org/the-coming-of-age-of-directed-energy-weapons-and-the-red-sea-crisis>
- Johnson, B., Ansley, J., Hakimipour, S., Buffin, K., Nguyen, L., Couture, V., & Gonaduwege, E. (2023). Counter directed energy weapons and the defense of naval unmanned aerial vehicles. *Journal of Directed Energy*, 7(2), 163–182. [https://nps.edu/documents/10180/142489929/JDE\\_7-2\\_Johnson.pdf](https://nps.edu/documents/10180/142489929/JDE_7-2_Johnson.pdf)
- Kim, E. S. (2024). *Propagation of beams carrying orbital angular momentum for directed energy weapons and free-space optical communications* [Master's thesis, Naval Postgraduate School]. <https://hdl.handle.net/10945/73160>
- Lazov, L., Teirumnieks, E., & Ghalot, R. S. (2021). Applications of laser technology in the Army. *Journal of Defense Management*, 11(4), Article 210. <http://dx.doi.org/10.35248/2167-0374.21.11.210>
- Middle East Council on Global Affairs. (2024). *Iran's missile and drone program: Disrupting U.S. aerial hegemony*. [https://mecouncil.org/wp-content/uploads/2024/07/ME-Council\\_Issue-Brief-Iranian-Drones-Final-\\_WEB.pdf](https://mecouncil.org/wp-content/uploads/2024/07/ME-Council_Issue-Brief-Iranian-Drones-Final-_WEB.pdf)
- Rassler, D. (2024). Going the distance: The emergence of long-range stand-off terrorism. *CTC Sentinel*, 17(2), 1–9. <https://ctc.westpoint.edu>
- Reeves, M., Haanaes, K., & Sinha, J. (2015). *Your strategy needs a strategy: How to choose and execute the right approach*. Harvard Business Review Press.
- Sabatini, R., Richardson, M. A., Gardi, A., & Ramasamy, S. (2015). Airborne laser sensors and integrated systems. *Progress in Aerospace Sciences*, 79, 15–63. <https://doi.org/10.1016/j.paerosci.2015.07.002>
- Sayler, K. M., DiMascio, J., Feickert, A., & O'Rourke, R. (2024). *Department of Defense directed energy weapons: Background and issues for Congress* (CRS Report No. R46925). Congressional Research Service. <https://crsreports.congress.gov/product/pdf/R/R46925>



- Shi, L., Pei, Y., Yun, Q., & Ge, Y. (2023). Agent-based effectiveness evaluation method and impact analysis of airborne laser weapon system in cooperative combat. *Chinese Journal of Aeronautics*, 36(4), 442–454. <https://doi.org/10.1016/j.cja.2022.11.006>
- Sylvester, J. E. (2014). *Power systems and energy storage modeling for directed energy weapons* [Master's thesis, Naval Postgraduate School]. <https://hdl.handle.net/10945/42734>
- Weinberg, G. V. (2023). *Directed energy system performance prediction*. Artech House.



# Embedded Capability Development: A Case Study in Rapid Missile Prototyping and Transition

**Maj Dillon T. Pierce, PhD**—is an active-duty Marine Corps officer serving as Chief Engineer for the Marine Corps Warfighting Laboratory Rapid Capabilities Office and as a Research Assistant Professor in the Space Systems Academic Group at the Naval Postgraduate School. He completed his doctoral research through the Marine Corps PhD Program—Technical (PHDP-T), where his dissertation focused on agile, low-cost missile development and flight validation. His current work centers on rapid capability development, missile prototyping, and transition-oriented defense innovation, which provides the basis for this paper.

**Mr. Christopher Reasonover**—is a government civilian detailed from Naval Surface Warfare Center Dahlgren Division and currently serving as the Precision Fire Control Program Officer within Code 352 (Kinetic Weapons) in the Office of Naval Research. His work spans the development, technical oversight, and transition of precision fires technologies relevant to future naval warfighting capability. He has an extensive background in radar and missile development, informed by 25 years of engineering experience at Dahlgren supporting advanced weapons and fire-control efforts. He holds a BS in Aerospace Engineering from Auburn University.

## Abstract

This paper presents embedded capability development as a practical execution approach for integrating requirements refinement, technical maturation, and transition preparation within a single collaborative government-industry development effort. In this model, a uniformed Service member is embedded within the technical team and serves both as a hands-on developer and as the principal government lead for the day-to-day refinement of evolving operational and system-level requirements. The approach addresses two persistent problems in defense capability development: promising science and technology efforts that fail to mature into transition-ready capabilities, and formal development processes that separate requirements, development, resourcing, and transition into sequential activities that delay delivery to the warfighter. The paper examines the model through an ongoing case study in low-cost tactical missile development and describes the organizational structure, collaborative industry engagement, and enabling tools used to implement it, including a government reference architecture and government reference design. The case suggests that embedded capability development can provide a practical complement to recent strategic-level reform guidance by showing how a program can organize itself to integrate technical learning, evolving requirements, and transition preparation within one collaborative development effort.

## Introduction

This paper examines an embedded capability development approach in which a uniformed Service member operates directly within a collaborative government-industry technical development effort while serving both as a hands-on developer and as the principal government lead for the day-to-day refinement of evolving operational and system-level requirements. Rather than treating requirements, development, and transition as sequential activities, the approach integrates them within a single iterative effort so that design decisions, prototype maturation, and requirement refinement inform one another continuously.

The approach is intended to address two persistent challenges in defense capability development. First, promising science and technology (S&T) efforts often fail to mature into fielded capability because the broader conditions required for transition are not established alongside technical progress. Second, formal capability development processes have historically struggled to adapt once technical and operational realities begin to evolve, particularly when requirements, resourcing, and acquisition activity remain too disconnected from ongoing technical learning. Embedded capability development is presented here as one



practical execution model for efforts that require meaningful technical development, involve transition complexity, and benefit from close collaboration across government and industry.

The paper first reviews the limitations of current capability development approaches, then introduces the model, and finally examines its implementation in an ongoing low-cost tactical missile development effort.

## **Background: Limitations of Current Capability Development Approaches**

Current military capability development approaches continue to exhibit two recurring failure patterns: promising technical efforts that do not transition beyond prototype demonstration, and formal development pathways that struggle to deliver integrated capability at the speed demanded by modern conflict.

### **The S&T Valley of Death**

A longstanding observation in defense innovation is the repeated failure of promising capabilities to cross the “Valley of Death,” the gap between technology development and acquisition. As the Government Accountability Office (GAO) and other Department of Defense (DoD) observers have noted, this gap is driven in part by a mismatch in maturation expectations: acquisition organizations typically require a higher degree of technical and programmatic maturity than S&T organizations are funded or incentivized to provide (DoD, 2007; GAO, 2015). The problem is compounded by the fact that successful transition depends not only on technical performance, but also on requirements alignment, integration planning, resourcing, sustainment, and a viable acquisition pathway. Because these functions are often distributed across separate organizations, transition frequently breaks down even when the underlying technology shows promise (Murray, 2023). In practice, promising prototypes may demonstrate feasibility, and even support limited operational employment, while still failing to establish the broader conditions necessary for acquisition and sustained fielding. The result is that potentially useful capabilities often do not mature into scalable, integrated systems that can be procured in quantity and employed broadly across the force.

### **Big “A” Capability Development and Current Reform**

A second limitation lies in the DoD’s formal capability development enterprise. “Big A” acquisition traditionally referred to the interaction of the Joint Capabilities Integration and Development System (JCIDS), Planning, Programming, Budgeting, and Execution (PPBE), and the Defense Acquisition System (DAS), which together linked requirements definition, resourcing, and acquisition execution in order to move a capability from identified need to funded program and fielded system (Lofgren, 2019). Although intended to provide discipline and accountability, this construct often operated more in sequence than in parallel. The result was a process that could be orderly, but slow to adapt when capability definition, technical maturation, funding alignment, and transition planning needed to inform one another during development.

That historical construct is now being actively reworked. In November 2025, the Department directed major reform of the joint requirements process, including establishment of the Requirements and Resourcing Alignment Board (RRAB), the Mission Engineering and Integration Activity (MEIA), and the Joint Acceleration Reserve (JAR), while also directing removal of references to JCIDS from relevant policy and associated guidance. In parallel, the Department redesignated the DAS as the Warfighting Acquisition System and issued broader acquisition transformation guidance oriented toward speed, adaptation, and delivery of operational capability (Office of the Secretary of War, 2025a, 2025b; Office of the Under Secretary of War for Acquisition & Sustainment, 2025). At the strategic level, these actions acknowledge the same underlying problem addressed in this paper: capability development processes that are too sequential, too administratively rigid, and too disconnected from iterative



technical learning. Even so, enterprise-level reform does not by itself specify how an individual development effort should be organized to operate differently in practice. That execution problem is where the embedded capability development model is intended to contribute.

## **The Embedded Capability Development Model**

Embedded capability development is proposed here as a practical execution model for programs that must mature technical capability, refine requirements, and prepare for transition in parallel. As shown in Figure 1, the model occupies a middle ground between rapid, prototype-driven experimentation and formal capability development processes. Its distinguishing feature is the concentration of technical development, accountable requirements refinement, and transition alignment within a single continuously evolving effort rather than across sequential organizational handoffs.

***Figure 7. Conceptual Depiction of the Embedded Capability Development Model***

Early stakeholder engagement, iterative prototyping, digital engineering, and government-industry collaboration are all familiar to the defense development community (GAO, 2023, 2025; Modigliani et al., 2020). What distinguishes the present model is their integration within a single execution construct centered on an embedded government developer who is not merely a technical participant, but the principal government lead for day-to-day refinement of evolving operational and system-level requirements. In this construct, that individual combines responsibilities that are often distributed across separate actors: direct technical participation, accountable refinement of requirements, continuous alignment with transition-relevant stakeholders, and stewardship of a retained government baseline. This changes not only how information moves across the effort, but when, where, and by whom critical development decisions are made.

### **Organizational Structure**

The organizational structure of the model is central to its function. For the model to work, the embedded developer must be operationally aligned within the chain of command responsible for capability development and requirements ownership while also operating inside a technical environment that supports direct design, build, test, and analysis. In this arrangement, formal authority over operational and system-level requirements remains with the appropriate capability development chain, but the embedded developer serves as the day-to-





day integrating point through which technical execution, evolving requirements, and transition considerations remain connected. This does not consolidate every institutional function into one person or office. Rather, it preserves the formal requirements authority of the chain of command while establishing a clear point of continuity between technical development and the organizations responsible for requirement refinement, transition, and broader capability development.

### **Enabling Functionality**

A second enabling feature of the model is the way it advances technical development, requirements refinement, and transition preparation together within the same effort. Rather than treating these as separate downstream activities, the model allows requirements to evolve alongside prototype evidence, keeps technical work tied to acquisition and integration realities, and generates the artifacts needed to support later transition and fielding. The result is not merely a working technical capability, but a more transition-ready one.

### **Core Drivers of Execution**

In practice, the effectiveness of the model depends on several reinforcing features that shape execution.

1. **Continuous technical participation by the embedded developer.** The embedded developer does not operate solely as a sponsor, coordinator, or oversight actor. The role carries direct involvement in technical development, allowing requirement refinement to remain informed by real design, build, test, and analysis activity rather than by intermittent review alone.
2. **Shared technical artifacts that preserve learning and enable iteration.** The effort maintains common technical artifacts that capture architecture, design knowledge, interfaces, verification logic, and implementation lessons. These artifacts provide a durable basis for collaboration, reduce learning loss across organizational boundaries, and allow the evolving baseline to be refined with greater continuity.
3. **Iterative requirements refinement informed by evidence.** Requirements are not treated as fixed end-state constraints established before meaningful technical learning occurs. Instead, they are refined through design, build, test, and trade-space evidence so that operational need, technical feasibility, manufacturability, and transition relevance can inform one another before the baseline hardens.
4. **Early industry collaboration against an evolving baseline.** Industry is engaged early enough to shape the effort rather than merely receive it. This allows manufacturing realities, implementation constraints, subsystem tradeoffs, and alternative solution paths to inform the baseline while design freedom is still high.
5. **Persistent alignment with transition, test, and operational stakeholders.** Transition, test, and operational stakeholders are engaged while the effort is still learning. This ensures that downstream constraints and opportunities influence development early, reducing the likelihood that promising technical efforts later stall because key realities were surfaced too late.

Taken together, these drivers form the practical execution logic of the model. They explain why embedded capability development is more than a coordination construct. It is a way of organizing a program so that technical learning, evolving requirements, and transition preparation reinforce one another during execution rather than converging only after the fact.



## Case Study: Low-Cost Tactical Missile Development

With the model established conceptually, this paper now turns to an ongoing case study in low-cost tactical missile development. Tactical missile development provides a useful test case because it combines high technical complexity, significant transition requirements, limited commercial sustainment incentives, and strong operational pressure for both speed and scale. These characteristics make it well suited for examining an approach intended to integrate requirement refinement, prototype maturation, and transition preparation within a single effort. If the model can add value in a domain as technically demanding and transition-heavy as tactical missile development, that strengthens the case for its applicability to other development-intensive military capabilities.

In the present case, the effort did not begin from a blank sheet. It built on a retained government technical foundation that had already established an initial system concept, generated flight-tested evidence, and captured early design and interface knowledge. That starting point shaped how the model was implemented in this program. It provided the initial basis from which later execution tools, including the government reference architecture and government reference design, could be developed and used to support collaborative refinement, technical maturation, and transition preparation. The discussion therefore begins with that baseline and the organizational context that enabled it before turning to the specific mechanisms through which the current effort is being executed.

### Government-Owned Technical Baseline

A defining feature of the present case is that it began with a government-owned technical baseline established before the current program construct was formed. That baseline originated in a government-sponsored research effort conducted at the Naval Postgraduate School (NPS). Executed as the primary author's doctoral dissertation, the effort examined lean, iterative approaches to the development of low-cost guided missile systems for future government capability development (Pierce, 2025). Although academically executed, the research was operationally motivated, sponsor-funded, and structured to generate reusable technical artifacts and insight for subsequent government efforts.

To establish that system-level missile baseline, the research employed an end-to-end design-build-test methodology that included the design, fabrication, and flight testing of a low-cost guided missile prototype. The purpose of the effort was not to produce a fielded weapon or acquisition-ready system, but to serve as an executable learning platform through which guidance, control, integration, and operational assumptions could be explored empirically. The resulting prototype, shown in Figure 2, provided a flight-testable starting point from which the later government-owned design baseline evolved. Data and observations from flight experiments informed the development of a government-owned reference architecture, initial interface assumptions, and an early understanding of system-level requirements. Retained as government-owned research outputs, these artifacts established a modular technical baseline structured to support experimentation, risk reduction, and requirements refinement independent of any specific vendor implementation or production pathway. As such, the effort provided both the technical foundation for the current case study and the basis for a follow-on organizational arrangement capable of carrying that baseline into continued collaborative development.





**Figure 8. Exterior View of the Flight-Test Prototype Developed During the Prior Government-Sponsored Research Effort that Established the Initial Government-Owned Technical Baseline for the Current Effort**

### **Implemented Organizational Structure**

Building on that baseline, the implemented organizational structure for this case study is shown in Figure 3 and reflects a dual-hat arrangement intended to connect formal capability development authority with a hands-on technical development environment. This structure is enabled in part by the flexibility of the Marine Corps' PhD Program–Technical (PHDP-T), which allows technically trained officers to be assigned dynamically to high-impact roles aligned to Service needs (HQMC CD&I, 2022). In this case, the embedded capability developer remains operationally aligned to Headquarters Marine Corps Combat Development and Integration (CD&I), reporting through Marine Corps Science and Technology / Rapid Capabilities Office leadership and coordinating across the relevant capability development organizations, including operational analysis, requirements, and transition stakeholders. More broadly, this arrangement is consistent with the Marine Corps' movement toward tighter integration of requirements, acquisition, and capability development functions under the newly established PAE-MC construct.

At the same time, technical execution is conducted at NPS, where faculty appointment and research access provide the facilities, technical ecosystem, and collaborative environment necessary to support hands-on design, build, test, and analysis. Additional coordination across NPS organizations further supports research direction, innovation activity, and technical execution. Within this arrangement, the embedded capability developer serves as the connective element between operational authority, technical development, transition coordination, and external collaboration, enabling rapid development cycles while preserving alignment with broader Service capability priorities. This organizational arrangement, however, is only one enabling condition of the realized model. Its effectiveness also depends on sustained, early, and technically meaningful collaboration with industry during execution.

***Figure 9. Organizational Realization of the Embedded Capability Development Model in the Current Case Study, Showing Command Alignment, Technical Execution Environment, and Collaborative Coordination Pathways***

### **Collaborative Missile Development with Industry Partners**

Complementing the organizational structure, a second key feature of the realized model is close technical collaboration with industry early in development. Once the initial government baseline had been established and feasibility demonstrated against an operationally relevant need, the effort was positioned for follow-on development within a broader naval capability program. It now operates as a deliberate, funded S&T effort focused on refining the technical baseline, interface definitions, subsystem performance parameters, and manufacturability considerations. In this phase, Office of Naval Research (ONR) sponsorship provides the resources necessary to move beyond the initial feasibility effort and support sustained development alongside structured industry participation.

Within that framework, collaboration with industry is conducted through Cooperative Research and Development Agreements (CRADAs) at NPS, allowing government and industry participants to work in close technical coordination while preserving government ownership of the broader technical baseline. As shown in Figure 4, this is not a one-way handoff from government requirements to industry execution. Rather, the government reference architecture (GRA), government reference design (GRD), and industry prototype path inform one another throughout execution. This allows prototype development, technical risk reduction, and requirements refinement to proceed together rather than through pre-stipulated handoffs. In this way, industry prototypes both inform and are informed by the evolving government baseline, allowing technical tradeoffs, interface definitions, and system parameters to be refined collaboratively before requirements are fully fixed. The practical mechanisms through which this collaboration is carried out are the GRA and GRD, which provide the shared technical basis for iterative requirements refinement, trade-space exploration, and prototype maturation.



**Figure 10. Dual-Path Development Structure Used in the Current Effort, Showing the Interaction Among the GRA, GRD, and Industry Prototype Path Within the ONR-Sponsored Development Effort**

## **GRA**

The first of these enabling tools is the GRA, defined as an authoritative government source of information that guides and constrains architectures and solutions within a specific subject area (Office of the DoD Chief Information Officer, 2010; Office of the Under Secretary of Defense for Research and Engineering, 2020). In traditional defense development, requirements and related baselines are often established early within broader sequential processes, while the technical learning needed to understand their downstream design, cost, schedule, and interface implications emerges later through development and testing. In the present approach, the GRA functions as a living development tool, evolving in parallel with the GRD and prototype evidence so that requirements, interfaces, and verification logic can be refined iteratively as technical understanding matures.

As a concrete example of the present implementation, Figure 5 depicts a high-level abstraction of one missile reference architecture developed in MATLAB System Composer. At this level, the architecture represents a mission thread for the missile system and is used to understand both the operational concept and the downstream system components required to realize it. Early in development, the GRA serves as a tool for trade-space analysis across alternative mission threads and operating concepts, allowing the team to evaluate how different approaches to achieving an operational effect drive different technical solutions and requirement implications.

The architecture shown in Figure 5 represents a digital mission-load autonomous strike concept in which the missile is preloaded from the launcher with targeting data, receives no inflight target updates, carries its own terminal guidance capability onboard, and is not





networked to other systems. That architecture drives specific subsystem and interface implications, including the need for an onboard terminal seeker and a launcher-to-missile data interface. By contrast, an architecture based on continuous laser designation changes both the operational context and the downstream system requirements because it depends on an external designator throughout the missile's flight and therefore alters the required subsystem configuration. In this way, the GRA helps illuminate the broader trade space associated with closing the kill chain, including operational constraints, interface needs, and the likely cost and performance implications of different architectural choices. As the architecture and reference design mature together, the GRA evolves from a lower-fidelity trade-space tool into a higher-fidelity authoritative source for requirements, interfaces, and verification logic, while also supporting automated verification of digital government and industry prototype solutions. The result is greater understanding when design freedom is highest and more informed requirements and system decisions as the effort matures.




**Figure 11. High-Level Abstracted Example of a Missile GRA Illustrating One Mission-Thread Implementation**

Although the embedded approach is not dependent on any single modeling tool, the choice of tooling materially affects how effectively the reference architecture and reference design can function as living development artifacts. In the present effort, the modeling environment is used not only to capture architectural relationships, but also to connect those relationships to higher-fidelity digital models and simulations associated with individual subsystems and prototype implementations. This allows the architecture to serve as more than a documentation product. It becomes part of an integrated digital ecosystem through which alternative configurations can be explored, prototype behavior can be assessed, and evolving requirements and interfaces can be verified against more realistic technical evidence.

## GRD

While the GRA defines the structure of the solution space, the GRD provides an executable system baseline through which those architectural decisions can be implemented and tested. In the present effort, the GRD is a government-retained missile system design containing the artifacts needed to build, replicate, test, and refine the capability in practice. These artifacts include physical design files such as computer-aided design and printed circuit board layouts, flight software, and associated simulation capability. As a result, the GRD serves as a practical vehicle for technical risk reduction, low-cost knowledge capture, and preservation of design knowledge that might otherwise remain locked inside a vendor-owned implementation.

More importantly, the retained government baseline is not merely a technical convenience. It is a strategic enabler of continuity, competition, and transition. By preserving technical learning inside a government-owned reference design, the model reduces dependency on any single performer, improves the government's ability to make informed transition decisions, and creates a stronger foundation for future competition, adaptation, and scaling. In this respect, the GRA and GRD support not only engineering development, but also continuity, transition planning, and future competition. Because the GRD is matured collaboratively with industry rather than in isolation, manufacturing realities, subsystem tradeoffs, and implementation constraints can shape design evolution early. In this way, the GRD helps generate more informed requirements, reduces the likelihood of late-stage transition failure, and provides a shared technical baseline through which government and industry can jointly burn down risk while preserving future flexibility.

As a concrete example of this coupling, Figure 6 depicts an abstracted portion of the missile flight software contained within the GRD. In this case, the software is developed in Simulink, auto-generated into flight code, and linked to a higher-fidelity six-degree-of-freedom simulation environment. This provides a direct connection between the evolving GRA and the underlying system realization. Interfaces represented in the flight software can be checked against the subsystem-to-subsystem contracts defined in the architecture, while alternative modules or subsystem implementations can be assessed against the same architectural requirements. In this way, the GRA and GRD are developed together rather than separately: the architecture informs the design, while the evolving design and simulation evidence refine the architecture. The result is a more integrated environment in which architecture, software, subsystem verification, and prototype maturation can converge together as the system develops.




**Figure 12. Abstracted Portion of the Model-Based Flight Software Within the GRD, Illustrating the Digital Implementation Layer Through Which Architectural Interfaces and Subsystem Contracts Can Be Refined and Verified**

## Execution of the Embedded Capability Development Model

The organizational structure, collaborative relationships, and enabling tools described above do not by themselves explain how the model functions in practice. In this program, the embedded capability development model is realized through the developer's continuous involvement in technical design, requirements refinement, stakeholder coordination, and transition preparation. Rather than treating these activities as distinct phases or separate organizational handoffs, the embedded developer owns and integrates them within the same ongoing development effort. The following discussion highlights three practical dimensions of that execution: iterative requirements development, stakeholder integration, and technical coordination with industry.

## Iterative Requirements Development

A central feature of the model in practice is that it begins from an operational need rather than from a fully specified set of operational and system-level requirements. In the present effort, the objective is not to translate a fixed requirement set downward into design, but to allow requirements to emerge and mature alongside design, build, test, and analysis activity. This reflects a different starting logic. Instead of assuming that the key system requirements are already known in sufficient detail, the effort treats many of them as matters to be discovered through technical development, trade studies, and empirical evidence. In this way, the model does not reduce requirements discipline; it relocates it into the learning process, where performance, manufacturability, cost, schedule, and operational utility can inform one another before the baseline hardens.

Figures 7 and 8 illustrate the distinction between the legacy linear requirements logic and the approach used in the present effort. Figure 7 depicts the historical deliberate requirements process as a reference point for the type of sequencing this paper seeks to overcome, while Figure 8 shows the coupled refinement process used here. In the traditional model, an operational need is translated relatively early into formal requirements, often before the technical implications of those requirements are fully understood. In the present approach, the operational need remains the starting point, but the detailed requirements needed to satisfy it are allowed to develop in parallel with the evolving technical baseline. Design evidence, build-test feedback, and implementation realities therefore shape the resulting requirement set rather than merely test compliance against an early assumption. The result is a process that is less dependent on prematurely fixed system definitions and more grounded in demonstrated knowledge about what the system can do, what tradeoffs are required, and what level of performance is realistically achievable to satisfy the operational need.

A concrete manifestation of this approach in the present effort is the use of set-based analysis before formal requirement lock. The program begins with a high-level operational need rather than a fully stipulated system definition, and alternative solution paths are explored against factors such as technical feasibility, manufacturability, integration burden, market availability, and transition relevance. In parallel, focused prototype efforts with the GRD are used to retire selected technical risks and capture key knowledge points. Those results feed back into the evolving baseline, allowing requirements to emerge from demonstrated evidence and informed trade decisions rather than from early assumptions about what the final system should be.




**Figure 13. JCIDS Deliberate Requirements Process. Adapted from Force Development System User Guide (United States Marine Corps, 2018).**

***Figure 14. Iterative Requirements Refinement Process Used in the Present Effort, Showing How Operational Need, Build-Test Feedback, Manufacturing Feedback, and Design Feedback Inform Coupled Operational and System Requirements Captured in the Digital Environment***

### **Early Stakeholder Integration and Transition Alignment**

A second practical dimension of the model is the deliberate early engagement of stakeholders that are often brought in only after a prototype has already matured. In the present effort, the embedded capability developer serves as the integrating point through which those organizations are pulled into the development process while the baseline is still evolving. As depicted conceptually in Figure 9, this includes the S&T sponsor, represented by ONR; Service capability development stakeholders within Headquarters Marine Corps CD&I; potential transition stakeholders within Program Executive Office Land Systems; operational test stakeholders within the Marine Corps Operational Test and Evaluation Activity; the technical execution environment at NPS; supporting technical capacity from University Affiliated Research Centers, Federally Funded Research and Development Centers, and Warfare Centers; and collaborating industry participants engaged through CRADAs.



***Figure 15. Conceptual Depiction of the Embedded Capability Developer as the Integrating Node for the Stakeholder Network Required to Mature the Capability Beyond a Standalone Prototype***

The purpose of this early integration is not simply to improve communication. It is to surface downstream realities early enough to shape the evolving baseline. Program office and systems engineering stakeholders can help ensure alignment with existing program trajectories, launcher and interface realities, and broader transition constraints. Operational test stakeholders can help identify opportunities to retire relevant test objectives during development rather than after the fact. Capability development and resourcing stakeholders can help keep the effort aligned with operational priorities and future programmatic considerations. In this way, the model pulls traditionally downstream considerations into the active development effort, reducing the likelihood that a technically promising prototype later stalls because key constraints, interfaces, or transition expectations were discovered too late.

### **Technical Collaboration with Industry Partners**

A third practical dimension of the embedded capability development model is the manner in which technical coordination with industry is carried out during execution. In this program, the embedded capability developer does not engage industry primarily through episodic reviews or formal handoff points, but through sustained day-to-day interaction within a shared technical development rhythm. This includes common access to relevant development artifacts and environments, such as software repositories, model-based design products, computer-aided design files, and interface definitions, together with recurring sprint-style engagements used to identify near-term objectives, technical blockers, and risk-burndown priorities.

What distinguishes this arrangement is that the embedded capability developer participates directly in low-level technical development while simultaneously interpreting those decisions against evolving operational requirements, transition constraints, and broader stakeholder priorities. This interaction is grounded not merely in recurring coordination, but in shared access to the evolving technical baseline, allowing discussion, trade studies, and implementation changes to occur against the same architecture, design, and interface context. Equally important, the collaboration framework reduces many of the practical administrative barriers that often slow government-industry interaction, including access to development environments, site access, and security-related coordination. The result is a form of collaboration that compresses technical, requirements, and transition feedback loops into the same development rhythm, reducing translation loss across organizational boundaries and





allowing government and industry to work side by side on technical maturation while preserving alignment with the larger capability development effort.

## **Early Observations from Implementation**

Although the effort remains ongoing, several practical observations can already be drawn from its implementation. These do not yet constitute a full assessment of transition outcomes, but they do show how the model is changing the development process in practice. Two examples are particularly illustrative: tighter coupling of architecture, design, and verification within the government baseline, and the use of set-based analysis before formal requirement lock.

One concrete example of execution impact in the present effort has been the maturation of the GRD's software and simulation environment through close collaboration with an industry partner. Earlier versions of the government baseline relied on a more monolithic flight software structure that was only loosely connected to the higher-fidelity simulation environment. Through collaborative development, that baseline was reworked so that model-based flight software modules could be auto-generated from simulation and exercised within the evolving implementation environment. This changed more than software development efficiency. It created a tighter connection between the GRD and the GRA, allowing digital interfaces, latency assumptions, navigation error parameters, and other subsystem-level considerations to be evaluated against a more realistic system realization earlier in development. In practical terms, this improved the quality of the government baseline itself by making later modular verification more credible and by allowing both government and industry implementations to be assessed against the same evolving architectural contracts and program objectives.

A second example is the use of set-based analysis before formal requirement lock. Rather than beginning with a fully stipulated system definition, the program starts from a high-level operational need and explores multiple feasible solution paths against factors such as technical feasibility, manufacturability, integration burden, market availability, and transition relevance. In parallel, focused prototype work with the GRD is used to retire selected technical risks and capture key knowledge points. The practical effect has been to shift requirement development away from early assumption and toward evidence-based convergence. Instead of treating requirements as fixed inputs to design, the present effort has allowed them to emerge from demonstrated trade studies, prototype results, and implementation realities. This has improved the technical grounding of the evolving requirement set while preserving flexibility during the early stages of development.

Taken together, these examples illustrate the practical value of the embedded approach in its current stage of execution. The benefit is not yet a completed transition outcome, but a stronger development baseline: one in which architecture, design, requirements, and transition considerations remain coupled while the system is still being defined. That integrated baseline should improve the quality of later transition decisions and reduce the likelihood of avoidable rework as the effort matures.

## **Conditions for Applicability**

The embedded capability development model is not a universal replacement for traditional acquisition, nor is it the only viable approach for accelerating capability delivery. It is best understood as one execution model within a broader set of emerging efforts to increase speed, flexibility, and operational relevance in defense development. Its greatest utility lies in a specific class of problem: capabilities that require meaningful technical maturation, involve nontrivial integration and transition complexity, and cannot be responsibly defined or acquired through either pure commercial purchase or a simple compete-then-commit logic at the outset.



First, the model is best suited to development-intensive capabilities that require genuine technical maturation and military integration but remain bounded enough to be advanced by a relatively agile technical effort. It is less relevant for pure commercial-off-the-shelf procurement, where the government can function primarily as a buyer, and less suited to very large, capital-intensive programs that require long-horizon program structures and extensive formal oversight from inception. Its natural fit lies between those extremes: efforts that must still be engineered, tested, and integrated, but that benefit from rapid iteration and tighter coupling between technical learning and capability definition.

Second, the model is most useful when the operational need is clear but the technical solution remains meaningfully uncertain. In such cases, architecture, subsystem selection, implementation constraints, manufacturability, and transition considerations must be learned together rather than assumed in advance. This distinguishes the model from approaches in which industry can be asked to build first against a relatively stable need and then compete for later production. Where the central challenge is not simply buying or selecting among candidate solutions, but discovering what the right requirement set should be as the system matures, embedded capability development offers a more appropriate mechanism.

Third, the model depends on an appropriately positioned government lead. The embedded developer must combine technical fluency, operational understanding, and sufficient institutional access to connect technical execution with broader capability development priorities. The role is especially well suited to personnel who can operate credibly across both engineering and operational communities while remaining nested within the chain of command responsible for capability development and formal requirements ownership. Without that alignment, the model risks losing the continuity that gives it value.

Finally, the model requires both a viable development environment and a plausible transition pathway. It depends on access to an environment capable of supporting real design, build, test, and analysis activity, together with the institutional relationships needed to carry a successful effort forward if the technical work proves promising. Without those conditions, the model risks becoming either a coordination construct without technical leverage or a prototype effort without a realistic path beyond demonstration.

## **Risks and Limits**

The model also carries important limits and risks. First, it depends heavily on personnel quality, institutional trust, and the availability of collaborative development authorities that may not exist in every program. If the embedded government lead lacks sufficient technical credibility, operational understanding, or institutional access, the model can quickly collapse into coordination without real developmental influence. Second, the model can create ambiguity if decision authority, government ownership boundaries, and transition expectations are not established clearly at the outset, particularly in a collaborative environment involving industry, sponsors, and potential transition stakeholders. There is also a risk of overextension if the embedded capability developer is treated as a substitute for the full requirements, acquisition, or test communities rather than as an integrating point across them. Third, while the model can reduce technical and transition risk, it does not eliminate the need for formal acquisition discipline, nor does it by itself solve the industrial-base, production-capacity, or long-term sustainment challenges associated with fielding capability at quantity.

A further limitation is that the strongest evidence for the model's value will often emerge only in retrospect. Because the present case remains ongoing, the observations offered here are necessarily preliminary. The paper therefore does not claim that the model has already proven universal success. It argues instead that the current effort demonstrates a plausible and



potentially valuable way to integrate development, requirements refinement, and transition preparation more closely than is typical in traditional approaches.

### Implications for Capability Development

Taken together, the observations, conditions, and risks discussed above suggest that embedded capability development is best understood as a practical way to organize collaborative government-industry development for capabilities that cannot simply be bought and cannot be responsibly specified in full detail at the outset. Its value lies in coupling technical learning, evolving requirements, industry realities, and transition considerations while the capability is still being defined. In that sense, it shifts the effort away from requirements development for its own sake and toward a more grounded process in which requirements emerge from design evidence, trade studies, implementation constraints, and operational need.

The result should be a better-informed capability baseline: one shaped not only by what is operationally desirable, but also by what is technically feasible, manufacturable, testable, and realistically transitionable. That does not replace formal acquisition. It helps set acquisition up better by improving the quality of the technical and programmatic baseline before later commitments harden.

### Recommendations

Based on the early execution of the model in the present case study, several preliminary recommendations emerge for applying embedded capability development more broadly.

1. **Apply the model selectively to development-intensive capabilities that cannot simply be bought and cannot be responsibly specified in full detail at the outset.** The model is most valuable where technical realization, operational utility, and transition considerations must be learned together rather than assumed in advance.
2. **Place the embedded developer inside both the technical development effort and the capability development chain responsible for requirements refinement.** The model depends on a lead who can participate directly in design, build, test, and analysis while maintaining continuity with the organizations responsible for evolving requirements, transition, and broader capability priorities.
3. **Use government-owned reference architectures and reference designs to support early collaborative development with industry.** Shared technical baselines provide a practical means of reducing risk, preserving government knowledge, improving requirements quality, and allowing government and industry to work against a common evolving baseline rather than through loosely coupled handoffs.
4. **Pull transition-relevant stakeholders and constraints left into development.** Early engagement with programmatic, systems engineering, test, and capability development stakeholders helps surface interface realities, transition constraints, and opportunities for risk retirement while the baseline is still evolving.

### Conclusion

This paper presented embedded capability development and examined its implementation in an ongoing low-cost tactical missile development effort. The model is intended for capability problems in which technical development, evolving requirements, and transition preparation must proceed together rather than sequentially. In the present case, that approach has been realized through a retained government baseline, a dual-hat organizational structure, close government-industry technical collaboration, and the use of evolving reference artifacts to support iterative development.



It is still too early to make definitive claims about ultimate transition success or the model's broader applicability. Even so, early evidence from the present effort suggests that embedded capability development provides a feasible way to organize collaborative government-industry development for capabilities that cannot simply be bought and cannot be responsibly defined in full detail at the outset. Its value lies in allowing technical learning, evolving requirements, industry realities, and transition considerations to shape one another while the capability is still being defined. For efforts of this kind, that may be its principal contribution: not as a substitute for acquisition, but as a practical way to produce a more coherent, better-informed, and more transition-ready baseline before later commitments harden.

## References

- DoD. (2007, August). *Tech transition report to Congress* [Report to Congress]. <https://www.dau.edu/sites/default/files/Migrated/CopDocuments/Tech%20Transition%20Report%20to%20Congress%20August%202007.pdf>
- DoD. (2020, September 9). *Department of Defense Directive 5000.01, The Defense Acquisition System* (Incorporating Change 1, July 28, 2022). <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/500001p.pdf>
- GAO. (2015, November). *Defense advanced research projects agency: Key factors drive transition of technologies, but better training and data dissemination can increase success* (GAO-16-5). <https://www.gao.gov/products/gao-16-5>
- GAO. (2016, November). *Weapon system requirements: Detailed systems engineering prior to product development positions programs for success* (GAO-17-77). <https://www.gao.gov/assets/gao-17-77.pdf>
- GAO. (2023, February 7). *Middle-tier defense acquisitions: Rapid prototyping and fielding requires changes to oversight and development approaches* (GAO-23-105008). <https://www.gao.gov/products/gao-23-105008>
- GAO. (2024, December 12). *DOD acquisition reform: Military departments should take steps to facilitate speed and innovation* (GAO-25-107003). <https://www.gao.gov/products/gao-25-107003>
- GAO. (2025, June 11). *Defense acquisition reform: Persistent challenges require new iterative approaches to delivering capability with speed* (GAO-25-108528). <https://www.gao.gov/products/gao-25-108528>
- Headquarters, United States Marine Corps, Combat Development and Integration. (2022, August 19). *Dynamic employment model decision paper: Marine Corps Doctor of Philosophy technical program dynamic employment model* [Decision paper].
- Lofgren, E. (2019, April 16). *What do we mean by big “A” and little “a” in acquisition?* Acquisition Talk. <https://acquisitiontalk.com/2019/04/what-do-we-mean-by-big-a-and-little-a-in-acquisition/>
- Modigliani, P., Ward, D., Lewis, T., & McGee, W. (2020, March). *Modernizing DoD requirements: Enabling speed, agility, and innovation*. MITRE Center for Technology and National Security. <https://www.mitre.org/sites/default/files/2021-11/prs-19-03715-2-modernizing-dod-requirements-enabling-speed-agility-and-innovation.pdf>
- Murray, K. (2023, March). *The Sea Services must build a new acquisition assembly line. Proceedings, 149(3/1,441)*. <https://www.usni.org/magazines/proceedings/2023/march/sea-services-must-build-new-acquisition-assembly-line>



- Office of the Department of Defense Chief Information Officer. (2010, June). *DoD reference architecture description*. Department of Defense.  
[https://dodcio.defense.gov/Portals/0/Documents/Ref\\_Archi\\_Description\\_Final\\_v1\\_18Jun10.pdf](https://dodcio.defense.gov/Portals/0/Documents/Ref_Archi_Description_Final_v1_18Jun10.pdf)
- Office of the Secretary of War. (2025a, November 7). *Reforming the joint requirements process to accelerate fielding of warfighting capabilities* [Memorandum].  
<https://media.defense.gov/2025/Nov/10/2003819442/-1/-1/1/REFORMING-THE-JOINT-REQUIREMENTS-PROCESS-TO-ACCELERATE-FIELDING-OF-WARFIGHTING-CAPABILITIES.PDF>
- Office of the Secretary of War. (2025b, November 7). *Transforming the Defense Acquisition System into the Warfighting Acquisition System to accelerate fielding of urgently needed capabilities to our warriors* [Memorandum].  
<https://media.defense.gov/2025/Nov/10/2003819439/-1/-1/1/TRANSFORMING-THE-DEFENSE-ACQUISITION-SYSTEM-INTO-THE-WARFIGHTING-ACQUISITION-SYSTEM-TO-ACCELERATE-FIELDING-OF-URGENTLY-NEEDED-CAPABILITIES-TO-OUR-WARRIORS.PDF>
- Office of the Under Secretary of Defense for Research and Engineering. (2020, November). *Mission engineering guide*. Department of Defense. [https://ac.cto.mil/wp-content/uploads/2020/12/MEG-v40\\_20201130\\_shm.pdf](https://ac.cto.mil/wp-content/uploads/2020/12/MEG-v40_20201130_shm.pdf)
- Office of the Under Secretary of War for Acquisition & Sustainment. (2025, November 10). *Acquisition transformation strategy: Rebuilding the arsenal of freedom* [Report].  
<https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/ACQUISITION-TRANSFORMATION-STRATEGY.PDF>
- Pierce, D. T. (2025). *Capability-cost inversion: An iterative framework for the development and flight-test validation of low-cost guided missiles* [Doctoral dissertation, Naval Postgraduate School, limited distribution].
- United States Marine Corps. (2018, April). *Force development system user guide*. Headquarters, United States Marine Corps.





# **Cyber Digital Twin-Informed Zero Trust: A Synergistic Framework for Securing Operational Technology in Defense Logistics Infrastructure**

**Dr. Barry A. Humphrey**—is the Program Manager for the Defense Logistics Agency's Research and Development Directorate. His program, Logistics Technology Research, aims at improving key business functions like data analytics, database security, supply chain risk management, and supply chain security through advanced technologies such as AI, Machine Learning, Blockchain, Augmented Reality, Quantum Computing, and Cybersecurity efforts to improve Warfighter support based on DLA's strategic guidance. Dr. Humphrey holds a PhD in Information Systems and Communications from Robert Morris University, an MA in Information Technology Management, Business and Organizational Security Management from Webster University, and a BS in Molecular Biology from Auburn University at Montgomery. He also completed Wharton Executive Education's Leadership Program in AI and Analytics and the Chief Technology Officer Program. Dr. Humphrey is co-founder and Director of Innovative Learning and Educational Technology at Ready Force Cyber, a non-profit organization providing cybersecurity, AI training and hands-on skill development, and career navigation in various technology fields. Dr. Humphrey is a retired military officer where he served over 22 years in the Air Force and Army.

## **Abstract**

The convergence of Information Technology (IT) and Operational Technology (OT) has exposed critical infrastructure to cyber-physical threats that perimeter-based security was never designed to handle. The consequences extend beyond data loss or equipment malfunction: compromised OT systems directly degrade military readiness, endanger both warfighter and civilian lives, and create national security vulnerabilities near-peer adversaries are actively probing. Legacy OT environments—the systems governing logistics, utilities, and manufacturing across military supply chains—operate under assumptions about isolation and trust that no longer hold. This research presents a security framework that integrates Cyber Digital Twins (CDT), Artificial Intelligence and Machine Learning (AI/ML), and a Zero Trust Architecture (ZTA) framework to provide an integrated defensive capability for OT cybersecurity. The approach centers on a high-fidelity virtual replica of the OT environment, training AI/ML models to recognize both normal operational signatures and simulated attack signatures within that replica, using the resulting risk intelligence to drive dynamic ZTA framework enforcement. The concept of the operational signature, the distinctive behavioral fingerprint of a device, process, or communication patterns is central to this framework: the CDT establishes baseline signatures, AI/ML models detect deviations from those signatures, and the ZTA framework enforces containment when anomalous signatures are identified.

**Keywords:** Cyber Digital Twin, Zero Trust Architecture, Operational Technology, AI/ML anomaly detection, operational signatures, defense logistics, national security, Industrial Control Systems

## **Introduction**

Anyone who has spent time inside a defense logistics operation understands the tension: these environments run on connectivity now—warehouse automation, predictive maintenance, enterprise resource planning integration—but the underlying Operational Technology (OT) systems were built in an era when cybersecurity meant locking the server room door. The digital transformation of defense logistics has delivered real operational gains, but it has also wired mission-critical Industrial Control Systems (ICS) into networks where sophisticated adversaries are establishing persistent, clandestine presence.

The Defense Logistics Agency (DLA) manages an extensive portfolio of OT systems that enable warfighter readiness: warehouse automation, environmental controls for material storage, fuel distribution infrastructure, and manufacturing equipment across the defense industrial base. These systems increasingly depend on network connectivity. That connectivity,



in turn, converts what were once air-gapped assets into potential entry points for nation-state actors and advanced persistent threats targeting military logistics (Cybersecurity and Infrastructure Security Agency, 2024). The national security implications are unambiguous. A compromised facility control system does not merely create a facility management problem—it can render temperature-sensitive munitions, pharmaceuticals, or electronics inventory unusable, directly degrading the readiness of deployed forces and placing both military personnel and the civilians they protect at risk.

The DLA Strategic Plan 2025–2030, “DLA Transforms: A Call to Action,” confronts this reality head-on. The “Precision” imperative pushes for stronger digital interoperability and artificial intelligence (AI)/machine learning (ML)–powered decision advantage. The “Posture” imperative demands that the Agency illuminate and mitigate global supply chain risk to increase resiliency (DLA, 2025a). The Director’s intent is blunt: to succeed in an environment where logistics are contested by adversaries across all domains and all levels of war, DLA must think, act, and operate differently. The CDT-ZTA Framework presented here is one answer to that mandate.

This research addresses these challenges through a Cyber Digital Twin–Informed Zero Trust Architecture (CDT-ZTA) Framework that emerged from multiple Small Business Innovation Research Phase I proof-of-concept projects conducted under the DLA R&D Directorate. The core argument is that integrating CDT simulation, AI/ML-driven signature analysis, and ZTA framework enforcement produces defensive capabilities none of the three can deliver alone—a security posture that is proactive, adaptive, and resilient rather than reactive and brittle.

### **Operational Technology Challenges**

OT is the hardware and software that monitors and controls physical devices, processes, and events (Gartner, 2023). The critical distinction from Information Technology (IT) is the priority stack: where IT traditionally leads with confidentiality, OT leads with availability and safety. A crashed email server is an inconvenience. A crashed programmable logic controller (PLC) governing a physical process in a defense logistics facility is a safety event, a readiness event, and a national security event.

This priority difference creates security challenges that IT-centric teams routinely underestimate. ICS deployed 20 or 30 years ago were built for isolated operation using proprietary protocols. Nobody designed encryption, authentication, or logging because there was nothing to connect to. Many of these systems are still running—they are reliable, replacement costs are enormous, and the operational risk of upgrading a functioning system is hard to justify (Hassanzadeh et al., 2020). Industry analysis from 2026 confirms the pattern: many OT environments still run firmware predating modern cryptographic standards or no longer supported by manufacturers. Protection rather than replacement remains the prevailing strategy (B2B Cyber Security, 2026).

From a national security perspective, the risk calculus is fundamentally different from enterprise IT. When adversaries compromise OT systems in defense logistics, they do not merely steal data; they gain the ability to manipulate physical processes that sustain military operations. Fuel distribution disrupted during deployment. Warehouse environments degraded to destroy critical inventory. Manufacturing equipment sabotaged to introduce defects into the defense supply chain. These scenarios translate directly to degraded readiness, mission failure, and risk to the lives of service members and the civilian populations they defend. The old framing of OT vulnerability as primarily a safety concern understates the stakes: OT security in defense logistics is a national security imperative.



The convergence trend has accelerated as organizations chase operational efficiencies—connected to monitoring, predictive maintenance, and enterprise integration. But every new connection expands the attack surface. The Dragos 2026 OT/ICS Cybersecurity Report documented that organizations with comprehensive OT visibility detected and contained ransomware incidents in an average of 5 days, compared with the industry-wide average of 42 days (Dragos, 2026). The SANS Institute’s analysis put it plainly: adversaries in 2025 showed patience and process awareness, investing in understanding how industrial environments work rather than developing novel malware. They were confident that long-standing architectural weaknesses would continue to deliver results (SANS Institute, 2026).




**Figure 1. The OT Threat Landscape Illustrating How Adversary Capabilities Target Defense Logistics OT Systems, with Consequences that Cascade from Facility-Level Disruptions to National Security Impacts Affecting Readiness and Lives**

## Zero Trust Architecture Framework Principles

The ZTA framework, at its core, is a rejection of the assumption that network location confers trustworthiness. Kindervag’s (2010) “never trust, always verify” principle has matured into a full architectural framework codified in NIST SP 800-207 (Rose et al., 2020). The operating assumption is that adversaries may already be inside the network—an assumption that OT security teams, having watched the erosion of air gaps over the past decade, find entirely plausible. A systematic literature review covering 74 peer-reviewed articles from 2016–2025 confirmed that the ZTA framework has transitioned from theoretical paradigm to practical implementation, with authentication, authorization, and access control as the most consistently deployed components (Al-Sharif et al., 2025). A broader review analyzing over 1,700 articles documented growing attention to OT-specific implementations (Alqahtani & Almuhaideb, 2025).

The ZTA framework tenets that matter most for OT are continuous verification, least privilege access, microsegmentation, and comprehensive monitoring. Microsegmentation is particularly valuable because it limits the blast radius—when an adversary gets in, they hit a wall rather than an open floor plan. Segment the OT network into granular trust zones with explicit access controls between them, and you contain threats before they reach safety-critical systems. In signal processing terms, microsegmentation constrains the propagation of anomalous signatures: a compromised device’s abnormal communication signature is contained within its segment rather than radiating across the entire network.

Implementing the ZTA framework in OT is harder than in enterprise IT, and anyone who says otherwise has not tried it. OT systems use legacy protocols with no authentication

mechanisms. They have real-time constraints where adding verification latency can disrupt physical processes. The Cloud Security Alliance's 2024 guidance on Zero Trust for Critical Infrastructure acknowledged these realities, offering a tailored five-step roadmap for industrial environments (Cloud Security Alliance, 2024). The Department of Defense (DoD) has moved aggressively on this front. In July 2025, DTM 25-003 directed all DoD components to achieve minimum Target Level Zero Trust across all systems, including OT and control systems (DoD, 2025a). In November 2025, the DoD CIO published "Zero Trust for Operational Technology Activities and Outcomes," defining 84 capability outcomes for target-level and 21 for advanced-level Zero Trust specific to OT (DoD, 2025b). Deadlines are set at fiscal 2030 for target-level and fiscal 2033 for advanced-level, with ZTA Strategy 2.0 expected in early 2026 (DefenseScoop, 2025). In January 2026, the NSA published a Zero Trust Implementation Guideline Primer providing phased direction aligned with the DoD ZTA Framework (National Security Agency, 2026).

### **Cyber Digital Twins for Security**

Digital twin technology—creating virtual replicas of physical systems for simulation and analysis—is well established in manufacturing and engineering. Extending the concept to cybersecurity means building CDTs that replicate not just functional behavior but security characteristics, vulnerabilities, attack surfaces, and critically, the operational signatures of every device, process, and communication flow in the environment (Eckhart & Ekelhart, 2019). The term "signature" is used deliberately. Just as radar systems identify aircraft by their electromagnetic signatures and sonar systems classify vessels by their acoustic signatures, a CDT captures the behavioral signatures of an OT environment: the timing patterns of PLC polling cycles, the characteristic data payloads of industrial protocol requests, the process response curves of physical equipment under load, and the network traffic profiles that constitute normal operations. These signatures become the baseline against which all deviations are measured.

The field has matured considerably. An ACM Computing Surveys review aggregating 201 publications mapped digital twins for security operations against the NIST Cybersecurity Framework, concluding that the technology sits at a convergence point of technology, cybersecurity, and industrial requirements (Vielberth et al., 2025). Repetto (2026) pushed further, proposing architectures for Cybersecurity Digital Twins that bridge fragmented security operations across multi-ownership digital service chains. The World Economic Forum highlighted CDTs as a transformative cybersecurity technology, noting that AI/ML has already cut breach detection times by 33% and containment times by 43% in security operations centers, and that CDTs build on those gains by enabling SOC's to simulate attacks and optimize defenses in real time (World Economic Forum, 2025). Industry has followed: Trend Micro announced enterprise-grade CDT technology for proactive cybersecurity in July 2025, powered by NVIDIA computing, spanning on-premises to cloud and IT to OT (Trend Micro, 2025).

For OT security specifically, the CDT's value is deeply practical. CDTs let you run attack simulations, test defensive measures, and validate security configurations without risking production systems. El Hachem et al. (2025) demonstrated that traditional threat modeling struggles with the dynamic, interconnected nature of ICS and that CDTs fill that gap with operational-context-aware security analysis throughout the system life cycle. By 2026, CDTs as virtual replicas for testing, simulation, and resilience training have become a defining trend in OT security, with organizations implementing application-aware microsegmentation and deep device signature fingerprinting within twin environments as standard practice (B2B Cyber Security, 2026). The DLA R&D portfolio has invested directly in these capabilities, with funding supporting CDT projects for process analysis and smart warehouse modernization (DLA, 2025b).



## Artificial Intelligence and Machine Learning for Anomaly Detection

AI/ML has fundamentally changed what is possible in cybersecurity analysis—automated pattern recognition and anomaly detection at volumes no human team can match. In OT environments, where normal operations produce consistent, predictable patterns of network traffic and process behavior, AI/ML models can learn the signature of what “normal” looks like and flag deviations that may indicate attack or malfunction (Zhou et al., 2021). The concept of the signature is again central: AI/ML models are, at their core, signature recognition systems. They learn the multidimensional signature of normal OT operations—comprising network traffic signatures, process variable signatures, device communication signatures, and timing signatures—and identify when observed behavior deviates from those learned baselines.

A systematic review published in *Electronics* in January 2026, analyzing 89 studies from 2021–2025, categorized AI/ML-based ICS anomaly detection across five data types: network traffic, operational data, simulation data, hybrid data, and auxiliary data. The key finding: network data catches cyber-layer attack signatures (reconnaissance, denial-of-service, man-in-the-middle), while operational data catches physical-layer anomaly signatures (process disturbances, false data injection, stealth deviations; Kim et al., 2026). This directly validates the CDT-ZTA Framework’s approach of combining both data types for comprehensive signature coverage.

The data scarcity problem in OT is real—production ICS environments do not experience enough attacks to train effective models from operational data alone. Recent work on Semi-Supervised Generative Adversarial Networks addressed this by using adversarial training to generate synthetic operational data that supplements limited real-world examples (Processes, 2025). This parallels the way the CDT-ZTA Framework uses its CDT to generate synthetic attack signatures. Birihanu and Lendák (2025) advanced explainable anomaly detection using multivariate Gaussian models that identify anomalous correlations between sensors and actuators while providing interpretable results—addressing the growing demand for Explainable AI/ML in critical infrastructure. At industry scale, Fortinet’s 2025 report confirmed the payoff: organizations on integrated OT security platforms with AI/ML-driven detection achieved up to 93% fewer cyber incidents compared to flat networks (Fortinet, 2025). The synergy between AI/ML detection and CDTs is straightforward: the CDT provides a safe environment for generating the training data—the attack and normal signatures—that AI/ML models need, while models deployed to both CDT and production enable continuous signature comparison and validation (Xu et al., 2022).

## Literature Review

The current literature reveals a clear trajectory: CDT technology, AI/ML, and the ZTA framework have each matured individually to the point where their convergence for OT security is not just feasible but increasingly necessary. The systematic reviews published between 2025 and 2026 tell a consistent story across independent research streams. Al-Sharif et al. (2025) and Alqahtani and Almuhaideb (2025) document the ZTA framework’s transition from conceptual model to deployable architecture, with growing attention to OT-specific implementations that account for legacy protocols, real-time constraints, and availability requirements that enterprise IT deployments never face. Vielberth et al. (2025) mapped 201 publications on CDTs for security operations against the NIST Cybersecurity Framework, concluding that the technology has reached a maturity level where it functions as an operational convergence point rather than a research curiosity. Kim et al. (2026), analyzing 89 studies on AI/ML-based ICS anomaly detection, found that hybrid approaches combining network traffic signatures and operational data signatures—the approach the CDT-ZTA Framework employs—represent the most effective detection strategy for industrial environments.





What the literature has not yet delivered is a validated integration architecture that ties these three capabilities together into a unified defensive system purpose-built for defense logistics OT—one where signature analysis is the connective thread. That gap is significant because the individual technologies create dependencies on each other that the literature acknowledges but has not resolved in practice. AI/ML-driven anomaly detection in OT environments is constrained by a well-documented training data problem: production ICS environments do not generate enough attack signatures to train effective models, and generating synthetic attack signatures on live systems risks operational disruption (Kim et al., 2026; Processes, 2025). CDTs solve that problem by providing safe environments for attack signature simulation and labeled dataset generation, but without an enforcement mechanism, detection insights remain informational rather than operational. ZTA framework microsegmentation provides that enforcement, but static segmentation policies degrade as environments change and adversaries adapt limitation that AI/ML-driven continuous risk assessment directly addresses.

Repetto (2026) argued that restricting security operations to individual administrative domains is inadequate against multi-step attacks that propagate across boundaries, while El Hachem et al. (2025) demonstrated that traditional threat modeling fails to capture the dynamic, interconnected nature of ICS environments. Literature, in other words, has identified the pieces and documented why each one alone is insufficient. What has been missing is the integration framework that leverages the CDT to generate and catalog operational and attack signatures, AI/ML to detect deviations from those signatures, and the ZTA framework to contain what the AI/ML inevitably misses, creating the kind of closed-loop, self-improving defensive architecture that the current threat landscape demands.

The threat intelligence literature reinforces the urgency from a national security standpoint. Dragos tracked 26 threat groups globally, including three newly identified, and reported a 49% year-over-year increase in ransomware groups hitting industrial organizations, affecting 3,300 organizations in 2025 (Dragos, 2026). KAMACITE conducted systematic reconnaissance of U.S. industrial devices throughout 2025, mapping entire control loops down to HMIs, variable frequency drives, and metering modules—essentially cataloging the operational signatures of American industrial infrastructure. ELECTRUM expanded beyond Ukrainian targets to European energy infrastructure with intent to disrupt operations. For defense logistics, translate those capabilities to ammunition production lines, fuel distribution networks, or warehouse control systems during a contested deployment, and the national security implications—degraded readiness, endangered military and civilian lives, strategic adversary advantage—are stark.

## **Methodology**

### **Research Design**

This research uses design science methodology. Design science focuses on architecting solutions to identified problems while generating knowledge that transfers beyond immediate implementation (Hevner et al., 2004). The CDT-ZTA Framework is the designed artifact; the problem is OT security in the defense logistics enterprise. The research team includes personnel from DLA's R&D department, cybersecurity division, and OT subject-matter experts, along with several trusted vendors currently under DLA R&D contract. This effort includes case-study analysis of framework deployment within DLA environments, allowing rigorous evaluation of framework effectiveness while capturing the contextual and organizational factors that determine whether a technically sound framework works in practice.



## Framework Architecture

The CDT-ZTA Framework integrates three components: the CDT environment, the AI/ML anomaly detection system, and the ZTA framework microsegmentation enforcement layer. The architecture defines distinct interfaces between them—data flows from physical systems to the CDT, from the CDT to AI/ML models, and from AI/ML risk assessments to ZTA framework policy enforcement. The concept of the operational signature is the connective tissue: the CDT captures signatures, AI/ML models learn and evaluate signatures, and the ZTA framework responds when anomalous signatures are detected.




**Figure 2. CDT-ZTA Framework Architecture Showing the Signature-Driven Data Flow from Physical OT Systems Through the CDT, AI/ML Engine, and ZTA Framework Enforcement, with the Continuous Improvement Loop that Drives Adaptive Defense.**

The CDT is a virtual replica of the target OT environment, including network topology, device emulations, and process simulations. It ingests real-time operational data from sensors and network-monitoring points in the physical environment, keeping synchronized with the actual system state. Critically, this data flow is one-directional—the CDT reflects current conditions while remaining isolated from the production network to prevent bidirectional propagation of risk. The CDT captures and catalogs the full spectrum of operational signatures: device communication signatures (protocol-specific traffic patterns, polling intervals, data payload structures), process signatures (physical response curves, setpoint trajectories, control loop behaviors), and network signatures (traffic volume baselines, flow directionality, service discovery patterns). These baseline signatures become the ground truth against which all future behavior is measured.

The AI/ML component runs multiple detection models trained on signature data generated within the CDT. Nominal operational data establishes behavioral baselines—the full set of normal signatures. Synthetic attack data generated through controlled threat simulation

within the CDT provides labeled examples of malicious signatures that supervised models need. The resulting models deploy to both the CDT (for continued refinement) and the production environment (for real-time monitoring). In signal processing terms, the AI/ML models function as matched filters: they are tuned to the known signatures of the environment and alert when incoming signals deviate from expected patterns beyond defined thresholds.

The ZTA framework microsegmentation component consumes continuous risk assessments from the AI/ML models and dynamically adjusts access control policies based on detected signature anomalies and predicted threat states. Software-defined networking enables granular enforcement down to individual devices. When the AI/ML system flags an anomalous signature from a specific device, microsegmentation policies automatically restrict that device's communication paths—containing the threat before lateral movement begins.




**Figure 3. The Signature Analysis Pipeline Showing How Operational Signatures Flow From Data Capture Through CDT Baseline, AI/ML Analysis, ZTA Framework Enforcement, and Back into the Learning Loop**

## Data Collection and Analysis

Data comes from multiple sources: network traffic captures, process variable logs, device state information, and security event data. Inside the CDT, comprehensive logging captures all system interactions during both normal operations and simulated attacks. This controlled environment produces labeled datasets with known ground truth—known normal signatures and known attack signatures—because the CDT generated both. Production data collection uses existing network monitoring infrastructure augmented with sensors positioned to capture OT-specific protocols and device communication signatures. Data anonymization and aggregation protect sensitive operational information while preserving the traffic patterns and behavioral signatures needed for analysis. Integration with DLA security operations centers provides additional threat intelligence and incident response context.

Analysis methods include statistical baseline characterization of operational signatures, AI/ML model performance evaluation (detection accuracy, false positive rates, detection latency, evasion resilience), and security effectiveness assessment through red team exercises. Red team evaluations pit skilled adversaries against systems protected by the framework, with results compared to baseline configurations without CDT-ZTA protections. The signature-centric approach enables precise measurement: how reliably do AI/ML models detect known attack signatures, how quickly do they identify novel signatures that deviate from baselines, and how effectively does the ZTA framework contain threats once anomalous signatures are flagged?

## Implementation Environment

Framework implementation is centered on a representative DLA facility OT environment.



In accordance with DLA CUI and security policies, specific system identifiers, facility names, and detailed equipment configurations are not disclosed in this paper. The test environment is representative of the Facility Related Control Systems (FRCS) common across DLA installations—systems that govern physical processes essential to maintaining the environmental, operational, and safety conditions required for mission-critical inventory storage and logistics operations. The U.S. Army Engineering and Support Center's Cybersecurity Systems Program has awarded contracts for logical and physical cybersecurity inventories of these systems at DLA locations, providing the foundational asset visibility that the CDT-ZTA Framework builds upon (DVIDS, 2024).

The test environment includes programmable logic controllers governing critical facility processes; supervisory control and monitoring systems coordinating operational parameters across multiple zones; local controllers managing individual subsystems; human-machine interfaces providing operators with real-time visibility into equipment status and conditions; and historian servers logging operational data for trend analysis and compliance reporting. These components are characteristic of the legacy OT architectures found across the DLA's logistics infrastructure—systems designed for reliability and availability that now require cybersecurity protections commensurate with their national security significance.

The CDT will be constructed on commercial virtualization platforms augmented with process simulation capabilities. Device emulations replicate the communication signatures and control logic of physical controllers. Network simulation accurately reproduces traffic signatures, including the periodic polling cycles, setpoint adjustments, and alarm conditions characteristic of the facility's OT operations. Process simulations model the physical dynamics of the controlled environment—so when an attack scenario manipulates a setpoint or disables a controller, the CDT reflects the downstream impact on process conditions and cascading equipment responses. This fidelity matters for signature analysis: the CDT must reproduce signatures with sufficient accuracy that AI/ML models trained within it perform reliably against real-world signatures in the production environment. An adversary who subtly adjusts a critical process setpoint may not trigger a conventional alarm, but the deviation in the operational process signature—detected by AI/ML models tuned to the CDT-established baseline—would flag the anomaly before the resulting conditions compromise mission-critical inventory or endanger personnel.

## **Results and Current Progress**

### **Phased Implementation Status**

The CDT-ZTA Framework is being implemented in phases, and it is important to present this work as it stands rather than as a finished product. Current efforts are concentrated on establishing the ZTA framework and microsegmentation foundation—the enforcement architecture that will eventually receive dynamic risk intelligence from the AI/ML and CDT components. This sequencing is deliberate: you cannot build adaptive, AI/ML-informed access control on top of a network that has no segmentation to adapt. The enforcement layer must exist before intelligence can drive it.

The current initiative centers on ZTA framework microsegmentation deployment. The work involves defining granular trust zones around individual OT subsystems—isolating controllers governing different physical processes from one another, separating supervisory management interfaces from local controllers, and enforcing explicit access policies between zones. Software-defined networking infrastructure enables policy enforcement at the device level, and initial segmentation rules are being validated against baseline traffic signatures to ensure that security policies do not inadvertently block legitimate operational communications.



The CDT and AI/ML anomaly detection components are planned for subsequent phases. The CDT environment will be constructed once the microsegmentation architecture is stable and validated, providing both a simulation platform for AI/ML model training and a pre-deployment testing environment for policy refinement. AI/ML models will then be trained within the CDT using both nominal operational signatures and synthetic attack signatures generated through controlled simulation, with the resulting detection capabilities layered onto the existing ZTA framework enforcement architecture.

### **Microsegmentation and ZTA Framework Progress**

Early results from the microsegmentation deployment are encouraging, though incomplete. Initial segmentation of the OT test environment has demonstrated the feasibility of device-level policy enforcement without disrupting operational processes. Communication between controllers and their supervisory systems has been restricted to only the specific data exchanges required for coordinated operations, eliminating the broad network access that previously allowed any device on the OT network to communicate with any other device. From a national security perspective, this is foundational: an adversary who compromises a single controller in a non-critical zone is now contained to that zone rather than gaining access to process-critical controllers, supervisory management interfaces, or historian servers—assets whose compromise could cascade into the kind of operational failures that destroy mission-critical inventory and degrade military readiness.

However, several challenges have emerged. Legacy controllers running outdated firmware have limited compatibility with modern SDN-based policy enforcement, necessitating workarounds that add management complexity. Some OT service discovery mechanisms rely on broadcast traffic that device-level segmentation restricts, requiring careful policy exceptions validated individually to avoid creating exploitable gaps. These are not insurmountable problems, but they illustrate why OT microsegmentation cannot be treated as a straightforward IT network exercise—every exception is a potential attack path, and every restriction is a potential operational disruption.

The broader industry data provides context for what mature microsegmentation should deliver. Fortinet's 2025 report found that organizations on integrated OT security platforms with device-level segmentation experienced up to 93% fewer cyber incidents than those on flat networks (Fortinet, 2025). Dragos documented that organizations with comprehensive OT visibility detected and contained ransomware incidents within 5 days, compared to the 42-day industry average (Dragos, 2026). These benchmarks represent the target state the CDT-ZTA implementation is building toward. The significance becomes concrete when mapped to documented threats. Dragos elevated VOLTZITE to Stage 2 of the ICS Cyber Kill Chain after observing the group manipulate engineering workstation software to extract configuration files and probe for operational conditions triggering process shutdowns (Dragos, 2026). In a DLA OT context, that translates to an adversary moving laterally from a compromised workstation to the controllers governing critical facility processes—exactly the kind of movement that device-level microsegmentation is designed to block.

### **Anticipated Operational Impact**

Because the project remains active in implementation, a final operational impact assessment is premature. What can be assessed are the anticipated impacts—both positive and negative—based on current progress, the test environment characteristics, and the broader evidence base.

If the framework performs as designed through full implementation, the anticipated benefits are substantial. Microsegmentation alone fundamentally changes the security posture by eliminating the flat network topology that allows unrestricted lateral movement. For a DLA





facility storing sensitive material—pharmaceuticals, electronics, munitions components—containment means the difference between a localized security incident and a facility-wide operational disruption that renders critical inventory unusable, with cascading consequences for military readiness and, ultimately, for the safety of military and civilian lives that depend on an unbroken logistics chain.

When the CDT and AI/ML layers are integrated, the anticipated capability uplift is significant. The CDT will enable pre-deployment validation of segmentation policy changes through signature analysis, eliminating trial-and-error on live systems. AI/ML-driven anomaly detection trained on OT-specific operational signatures should identify subtle attacks—a gradual process setpoint manipulation producing an anomalous process signature, an unauthorized polling pattern producing an anomalous network signature—that rule-based monitoring would miss entirely. The combination of rapid AI/ML signature deviation detection feeding dynamic microsegmentation enforcement should compress the window between intrusion and containment from hours or days to seconds.

The risks of getting this wrong are not abstract. A misconfigured microsegmentation policy that blocks a critical communication path between a controller, and its supervisory system does not just generate an alert can disrupt the physical process that the controller governs, with potential consequences for both inventory integrity and personnel safety. False positives from the AI/ML detection layer, once integrated, present a related risk. If AI/ML models are not adequately trained on the full range of legitimate operational signature seasonal variations, maintenance cycling, demand response events, equipment degradation patterns, they will flag normal operations as suspicious. There is also the risk that adversaries specifically target the framework's integration points, manipulating what the CDT sees to corrupt baseline signatures or trigger inappropriate ZTA framework responses—the kind of patient, process-aware adversary behavior that Dragos documented in KAMACITE's systematic mapping of U.S. industrial control loops throughout 2025 (Dragos, 2026).




**Figure 4. The Synergistic Integration Model Shows How CDT, AI/ML, and the ZTA Framework Components Depend On One Another. The CDT Generates Signatures for AI/ML Training, AI/ML Provides Risk Intelligence to the ZTA Framework, and the ZTA Framework Validates Policies Within the CDT—Creating a Closed-Loop, Signature-Driven Defensive Architecture.**

## Conclusions

This research is in early stages, and saying otherwise would misrepresent where the work stands. What can be said is that the CDT-ZTA Framework's foundational phase—establishing ZTA framework microsegmentation within a representative DLA OT environment—has demonstrated that device-level policy enforcement in legacy OT environments is achievable without compromising the availability that mission-critical facility control systems demand. That is a necessary first step; it is a promising one.

The early results validate the core architectural premise: build the enforcement layer first, then add signature-based intelligence on top of it. Microsegmentation is establishing the trust zone boundaries and access controls that will eventually receive dynamic risk assessments driven by AI/ML signature analysis from the CDT components. The challenges encountered so far—legacy firmware incompatibilities, OT broadcast dependencies, the operational complexity of managing device-level policy exceptions—are real but solvable and documenting them provides the implementation lessons that other DoD organizations will need as they pursue their own OT ZTA framework mandates under DTM 25-003 and the November 2025 OT-specific guidance.

The national security stakes elevate this work beyond technical exercise. The threat landscape is not waiting for the research to conclude. Dragos documented adversary groups actively mapping industrial control loop signatures across U.S. infrastructure throughout 2025, with KAMACITE conducting systematic reconnaissance and VOLTZITE escalating to Stage 2 of the ICS Cyber Kill Chain (Dragos, 2026). Ransomware groups targeting industrial organizations surged 49% year-over-year, impacting 3,300 organizations. For DLA facilities where OT systems are critical for mission success, these are not theoretical risks—they are active adversary capabilities being refined against the same classes of systems this framework is designed to protect. Compromised OT systems in defense logistics mean degraded readiness, endangered military and civilian lives, and a strategic advantage surrendered to adversaries. The CDT-ZTA Framework is designed to prevent that outcome.

The DLA Strategic Plan 2025–2030 calls for the Agency to think, act, and operate in new ways—and doing so responsibly means validating each phase before building on it, being transparent about what works and what does not, and maintaining the intellectual honesty to adjust course as results and threats evolve. The CDT-ZTA Framework, with its signature-centric approach to detection and containment, is on a trajectory to deliver the integrated defensive capability that DLA's OT environments need. The framework aligns with DTM 25-003, November 2025 OT-specific ZTA framework guidance, and the NSA's January 2026 Zero Trust Implementation Guideline Primer, positioning DLA to meet compliance deadlines while building capabilities that improve security rather than merely satisfying checklists. The start is promising. The work continues.

## Future Steps

The DLA R&D Directorate will continue this research through subsequent phases that expand the framework's capabilities as the technology landscape evolves. Near-term priorities include constructing the CDT of the OT test environment and initiating AI/ML model training using both nominal operational signatures and synthetic attack signatures generated within the CDT. The signature library, the catalog of known normal and known malicious behavioral fingerprints—will grow with each phase, becoming an increasingly powerful asset for detection and response.

Several emerging technology integrations merit focused R&D attention. Large Language Model (LLM) integration into the CDT ecosystem represents a near-horizon opportunity. Recent



work has demonstrated that LLMs can automate interpretation of complex security telemetry, assist real-time threat analysis, and generate human-readable explanations that make sophisticated signature detection outputs accessible to facility operators who are not data scientists (Ferraro et al., 2025). Post-quantum cryptographic protocols for securing OT communications channels will become critical as quantum computing matures. Federated learning approaches that enable AI/ML models to improve across multiple DLA installations without centralizing sensitive operational signature data address both the data scarcity problem and data sovereignty requirements. Autonomous response orchestration that compresses the decision loop between signature deviation detection and containment is another area under evaluation.

Scalability requires candid assessment. CDT demands significant computational resources and specialized expertise for accurate OT modeling. Organizations need to decide whether to build that capability in-house, leverage commercial platforms, or pursue a hybrid approach. The AI/ML component requires ongoing care: models trained today will degrade as adversaries evolve their attack signatures and systems change their operational signatures. Continuous retraining, integration of current threat intelligence, and model validation are operational requirements for the life of the system. Enterprise-scale microsegmentation introduces network management complexity addressable only through automation and orchestration.

Organizational change management remains essential. OT operations personnel need to understand how security controls affect their systems and have clear channels for identifying legitimate communications that policies might inadvertently block. Security personnel need to develop OT-specific expertise—investigating alerts in an industrial environment is fundamentally different from investigating alerts in an enterprise IT network. The DLA Strategic Plan's emphasis on organizational agility and mission-driven skill development supports these management requirements. Technical capability without organizational readiness is expensive shelfware.

As the DoD develops additional ZTA framework guidance for weapon systems and defense-critical infrastructure (DefenseScoop, 2025), extending the CDT-ZTA Framework to those domains—adapting the architecture to the unique operational signatures and constraints of weapons platforms and mission-critical infrastructure—is a natural and necessary evolution. Broader ambition is a framework that does not just keep pace with adversary evolution but anticipates it—one that learns new signatures, adapts its models, and extends to new DLA environments and technology domains as the defense logistics mission demands. Warehouse automation, fuel distribution infrastructure, and manufacturing equipment across the defense industrial base—these all present variations of the same fundamental challenge: legacy OT systems connected to modern networks, operating under availability and national security constraints that make traditional IT security approaches inadequate and sometimes dangerous.

## References

- Ahmad, Z., et al. (2025). A new era for digital twins: Progress and industry adoption. *International Journal of Computer Integrated Manufacturing*.  
<https://doi.org/10.1080/27525783.2025.2555877>
- Al-Sharif, Z. A., et al. (2025). A systematic literature review on the implementation and challenges of Zero Trust architecture across domains. *Sensors*, 25(19), 6118.  
<https://doi.org/10.3390/s25196118>
- Alqahtani, A., & Almuhaideb, A. (2025). *Comprehensive review of Zero Trust research landscape* [Reference details].



- B2B Cyber Security. (2026, February). *Digital twins for modern OT resilience*. <https://b2b-cyber-security.de/en/digitale-zwillinge-fuer-moderne-ot-resilienz/>
- Birihanu, E., & Lendák, I. (2025). Explainable correlation-based anomaly detection for Industrial Control Systems. *Frontiers in Artificial Intelligence*, 7, 1508821. <https://doi.org/10.3389/frai.2024.1508821>
- Cloud Security Alliance. (2024). *Zero Trust guidance for critical infrastructure*.
- Cybersecurity and Infrastructure Security Agency. (2024). *Industrial control systems cyber emergency response team year in review 2023*. U.S. Department of Homeland Security.
- Defense Logistics Agency. (2025a). *DLA transforms: A call to action—Strategic plan 2025–2030*. <https://www.dla.mil/Info/Strategic-Plan/>
- Defense Logistics Agency. (2025b). *Defense Logistics Agency FY2026 research, development, test & evaluation budget justification*. <https://comptroller.war.gov/>
- DefenseScoop. (2025, December 9). *Pentagon plans to publish Zero Trust strategy 2.0 in early 2026*. <https://defensescoop.com/2025/12/09/dod-zero-trust-strategy-2-0-expected-early-2026/>
- Dickson, P. (2025). *IDC assessment of digital twin technology for cybersecurity exposure management*. IDC.
- DoD. (2025a). *DTM 25-003: Implementing the DoD Zero Trust strategy*. <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dtm/DTM%2025-003.PDF>
- DoD. (2025b). *Zero Trust for operational technology activities and outcomes*. DoD CIO Zero Trust Portfolio Management Office. <https://dodcio.defense.gov/Portals/0/Documents/Library/ZTOperationalTechnologyActivitiesOutcomes.pdf>
- Dragos. (2026). *2026 OT/ICS cybersecurity report and year in review*. <https://www.dragos.com/ot-cybersecurity-year-in-review>
- DVIDS. (2024). Cybersecurity program providing vital services for National Guard, Defense Logistics Agency. <https://www.dvidshub.net/news/489056/>
- Eckhart, M., & Ekelhart, A. (2019). Digital twins for cyber-physical systems security: State of the art and outlook. In *Security and quality in cyber-physical systems engineering* (pp. 383–412). Springer.
- El Hachem, J., et al. (2025). Leveraging digital twins for advanced threat modeling in cyber-physical systems cybersecurity. *International Journal of Information Security*. <https://doi.org/10.1007/s10207-025-01043-x>
- El Ouedrhiri, Z., et al. (2025). An explainable Markov chain–machine learning sequential-aware anomaly detection framework for industrial IoT systems based on OPC UA. *Sensors*, 25(19), 6122.
- Ferraro, A., et al. (2025). Enabling cyber security education through digital twins and generative AI. *arXiv preprint*, arXiv:2507.17518.
- Fortinet. (2025). *2025 State of operational technology and cybersecurity report*.
- Gartner. (2023). *Information technology glossary: Operational technology*.
- Hassanzadeh, A., Burkett, R., & Xu, S. (2020). A cyber-physical testbed for detecting attacks on industrial control systems. *Proceedings of the 2020 ACM SIGSAC Conference on*



- Computer and Communications Security*, 2211–2213.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105.
- Kim, J., et al. (2026). AI-based anomaly detection in industrial control and cyber-physical systems: A data-type-oriented systematic review. *Electronics*, 15(1), 20. <https://doi.org/10.3390/electronics15010020>
- Kindervag, J. (2010). *Build security into your network's DNA: The Zero Trust network architecture*. Forrester Research.
- National Institute of Standards and Technology. (2023). *Guide to operational technology security* (NIST Special Publication 800-82 Rev. 3). U.S. Department of Commerce.
- National Security Agency. (2026). *Zero Trust implementation guideline primer* (U/OO/102936-26). [https://media.defense.gov/2026/Jan/08/2003852320/-1/-1/0/CTR\\_ZERO\\_TRUST\\_IMPLEMENTATION\\_GUIDELINE\\_PRIMER.PDF](https://media.defense.gov/2026/Jan/08/2003852320/-1/-1/0/CTR_ZERO_TRUST_IMPLEMENTATION_GUIDELINE_PRIMER.PDF)
- Processes. (2025). *Semi-supervised generative adversarial networks for synthetic OT data generation*. [Reference details].
- Repetto, M. (2026). Cybersecurity digital twins: Concept, blueprint, and challenges for multi-ownership digital service chains. *Journal of Information Security and Applications*, 96, 104299. <https://doi.org/10.1016/j.jisa.2025.104299>
- Resnick, J. (2025). *Zero Trust Portfolio Management Office commentary on OT requirements* [Reference details].
- Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero trust architecture* (NIST Special Publication 800-207). U.S. Department of Commerce.
- SANS Institute. (2026). *Top takeaways from the Dragos 2026 OT cybersecurity report*. <https://www.sans.org/blog/top-takeaways-from-the-dragos-ot-cybersecurity-report-2026>
- Stouffer, K., Pease, M., Tang, C., Zimmerman, T., Pillitteri, V., & Lightman, S. (2023). *Guide to operational technology security* (NIST Special Publication 800-82 Rev. 3). National Institute of Standards and Technology.
- Trend Micro. (2025, July 31). Trend Micro reinvents proactive security with digital twin technology. <https://newsroom.trendmicro.com/2025-07-31>
- Vielberth, M., et al. (2025). Digital twins in security operations: State of the art and future perspectives. *ACM Computing Surveys*. <https://doi.org/10.1145/3746279>
- World Economic Forum. (2025, March). How digital twin technology can enhance cybersecurity. <https://www.weforum.org/stories/2025/03/how-digital-twin-technology-can-enhance-cyber-security/>
- Xu, J., Wen, H., Li, M., & Chen, X. (2022). Digital twin-enabled anomaly detection for industrial control systems. *IEEE Internet of Things Journal*, 9(15), 13585–13595.
- Zhou, Y., Zhang, X., Li, Y., & Yang, Q. (2021). A survey on deep learning approaches for network intrusion detection systems in industrial control systems. *IEEE Access*, 9, 85379–85402.





THIS PAGE INTENTIONALLY LEFT BLANK



## PANEL 4. THE FRICTIONLESS FRONTIER: REFORMING ACQUISITION TO SECURE CUTTING-EDGE CAPABILITY

Wednesday, May 6, 2026, 1145 – 1300 ET (0845 - 1000 PT)

### Panel Summary:

Preserving military superiority requires an acquisition system that can successfully attract non-traditional innovators while maintaining the agility to manage complex contracts under "wartime" pressure. This panel addresses the systemic barriers preventing new commercial entrants from joining the defense industrial base and introduces proactive frameworks—from wargaming the contracting lifecycle to "re-risking" through contract-type conversion—designed to increase budget predictability and mission resilience. By aligning institutional incentives with warfighting outcomes, these research initiatives offer a roadmap for a more inclusive, flexible, and responsive defense marketplace.

**Chair: Mr. John Jolokai, Executive Director and Senior Contracting Officer (SCO), U.S. Army Contracting Command- Detroit Arsenal**

### Presenters:

**Exploring the Use Case of Wargaming Tools Across the Contracting Lifecycle – Lt Col Jamie Porchia, Assistant Professor, Naval Postgraduate School**

**Is the Military Attracting New Companies? Assessing First-Time Entrants into the Defense Market – Amanda Bresler, President, PW Communications**

**Buying Flexibility: Converting Cost-Plus to Firm-Fixed-Price to Strengthen Mission Resilience – Martin Jay Barlow, Senior Engineer Specialist, The Aerospace Corporation**



**Mr. John Jolokai**—is the Executive Director and Senior Contracting Officer (SCO) of the Army Contracting Command - Detroit Arsenal (ACCDTA). ACC-DTA ensures warfighting readiness by providing contracting and acquisition support for combat and tactical vehicle systems, deployment and Soldier support equipment, and intelligence and security related requirements. Mr. Jolokai leads over 700 civilian and military personnel across seven sites, executing more than 20,300 contracting actions annually, totaling \$10 billion in obligations.

Prior to becoming the ACC-DTA Executive Director, Mr. Jolokai served as the Deputy to Commander, U.S. Army Contracting Command, 409th Contracting Support Brigade (CSB) in Kaiserslautern, Germany. The 409th Contracting Support Brigade provides enhanced readiness and operational capacity for U.S.

and Partner Forces in the European Theater by delivering the power of Army contracting to leverage commercial capability and provide business advice to support sustaining, shaping, and decisive operations to maintain a strong Europe.

After starting with the TACOM Acquisition Center in 2001, Mr. Jolokai has held several key contracting roles, including Chief of the ACC-DTA Source Selection Center of Excellence, Tactical/MRAP Contracting Division Chief, Civilian Deputy Army Contracting Command-Detroit Arsenal ACC-DTA, and assignments with the Office of the Deputy Assistant Secretary of the Army (Procurement) and the Army's Next Generation Combat Vehicles Cross-Functional Team.

Mr. Jolokai holds a Master's in Global Leadership Management (2019), MBA (2008), and a Bachelor of Arts in Finance (2003). He is a member of the Army Acquisition Corps and is certified in Contracting. He is a graduate of the 2019 Defense Acquisition University Senior Service College.



# Exploring the Use Case of Wargaming Tools Across the Contracting Life Cycle

**Jamie Porchia, Lt Col, PhD**—is an Assistant Professor at the Naval Postgraduate School where she teaches enterprise innovation, strategic sourcing, and introductory contracting courses. She is an Air Force contracting officer with more than 17 years of experience across a variety of contracting specialties. She has multiple publications in the International Journal of Physical Distribution & Logistics Management, Supply Chain Management: An International Journal, and the Journal of Contract Management. Her primary research interests are supply chain mapping, drone supply chains, supply chain design and procurement.

**Michael Alexander, Capt, MSCM**—enlisted in the Oklahoma ANG 2014. While in the Guard studied Accounting at the University of Oklahoma and joined Air Force ROTC subsequently commissioning in May 2019. Capt. Alexander has been stationed at Seymour Johnson AFB, NC; Hanscom AFB, MA; and the Naval Postgraduate School. Previously, he was the sole warranted contracting officer supporting the \$6.8 billion ACAT IC FAB-T program, where he led and managed complex negotiations. He has served in multiple leadership roles—including contracting specialist, contracting officer, lead contracting officer, team lead, and flight commander—demonstrating excellence across the contracting life cycle, contingency and operational contracting, and system-level contracting.

## Abstract

From the receipt of a requirement to the close of a contract, there are numerous activities that occur throughout the contract life cycle. These activities are performed by contracting professionals who are faced with making business judgements that balance multiple tradeoffs such as awarding contracts with speed at a fair and reasonable price, meeting the mission's needs in a timely manner, mitigating risks, and adhering to laws and regulations. Nonetheless, contracting professionals are required to be not only technically proficient but professionally competent as well. To develop these requisite competencies, incorporating wargaming tools into the education and training fabric of Defense contracting offers a plausible path for helping contracting professionals attain these important skills. This exploratory and conceptual research focuses on how the use of wargaming tools can be integrated into the contracting enterprise. Through a review of extant literature, a framework is proposed for developing educational wargames and describing how a wargame's objectives can be crafted for contracting specific games. A notional use case is developed showing how a wargame can be developed to "battle test" an acquisition plan.

## Introduction

From the receipt of a requirement to the close of a contract, there are numerous activities that occur throughout a Defense contract's life cycle. These activities are performed by contracting professionals that are trained to a standard of competencies governed by the Department of Defense Contracting Competency Model (DoDCCM; DoD, 2024). Contracting professionals are faced with making business judgements that balance multiple tradeoffs such as awarding contracts with speed at a fair and reasonable price, meeting the mission's needs in a timely manner, mitigating risks, and adhering to laws and regulations (Federal Acquisition Regulation 1.602-1, 2026). Also noting, these tradeoffs are made within the context of a dynamically shifting acquisition environment.

Nonetheless, contracting professionals are required to be not only technically proficient but professionally competent as well. Professional competencies include creativity and innovation, critical thinking, problem solving, risk management and strategic thinking, to name a few (DoD, 2024). To develop these requisite competencies, incorporating wargaming tools into the education and training fabric of Defense contracting offers a plausible path for helping contracting professionals attain these important skills. Meticulously crafted wargames can



facilitate learning, training and testing of human judgement in a fail-safe environment (Lorusso, 2024), thus allowing contracting professionals to practice their decision-making in a low-stakes environment and before decisions must be made. Furthermore, wargames are tailorable to the specific needs of the sponsor, or to meet training objectives (Appleget et al., 2020) like those that could be derived from the DoDCCM's prescribed competencies. To explore the mechanism that could enable such a process, this research examines the following research questions:

RQ1: How can wargaming tools be integrated into contracting life cycle processes?

RQ2: How can wargaming be utilized as a training and preparedness tool for contracting professionals?

To address these research questions, this paper begins with a review of pertinent literature on the contracting life cycle, contracting competencies, and wargaming. After the literature review, the framework will be presented followed by an example application. The paper concludes with future research recommendations.

## **Literature Review**

### **Contract Life Cycle and Contracting Competencies**

The DoDCCM is the governing source for competency standards for Department of War (DoW) contracting professionals. The DoDCCM adopted the National Contract Management Association's (NCMA) Contract Management Standard™ (CMS) as its basis for developing a standardized contracting workforce (DoD, 2024). There are five components within the DoDCCM, of which the contract life cycle phases, competencies, and tasks are most applicable to this paper. The contract life cycle represents the three phases of a contract as it moves from pre-award to award, then subsequently post-award. Within each of the phases there are five domains that capture the breadth of actions performed by contract professionals (Commerce and Contract Management Institute, 2019). Each domain consists of several common contracting competencies that all contracting professionals within the DoW should understand along with accompanying job tasks that enable the achievement of the competencies. Figure 1 is the DoDCCM. The blue rectangles represent the contract life cycle phases, the green rectangles represent the domain, and the purple rectangles are the competencies. Figure 2 is an example of an extended Figure 1 that includes the tasks for the Develop Offer domain.






**Figure 1. DoDCCM (DoD, 2024)**



|                   |                                           |                                                                      |
|-------------------|-------------------------------------------|----------------------------------------------------------------------|
| 2.2 Develop Offer | 2.2.1 Plan Sales<br>[Seller Job Tasks]    | 2.2.1.1 Conduct Pre-Sales Activities                                 |
|                   |                                           | 2.2.1.1.1 Assess Customer Relationships                              |
|                   |                                           | 2.2.1.1.2 Develop Marketing Strategy                                 |
|                   |                                           | 2.2.1.1.3 Assess Competition                                         |
|                   | 2.2.1.2 Evaluate Solicitation             | 2.2.1.1.4 Determine Supply Chain Support                             |
|                   |                                           | 2.2.1.2.1 Request Clarification                                      |
|                   |                                           | 2.2.1.2.2 Propose Solicitation Changes                               |
|                   | 2.2.1.3 Conduct Offer / No-Offer Analysis |                                                                      |
|                   |                                           |                                                                      |
|                   | 2.2.2 Prepare Offer<br>[Seller Job Tasks] | 2.2.1.4 Finalize Sales Plan                                          |
|                   |                                           | 2.2.2.1 Execute Sales Plan                                           |
|                   |                                           | 2.2.2.2 Develop Execution Plan                                       |
|                   |                                           | 2.2.2.2.1 Understand Unique and Special Requirements                 |
|                   |                                           | 2.2.2.2.2 Assess Capability to Satisfy All Solicitation Requirements |
|                   |                                           | 2.2.2.3 Develop Risk Mitigation Plans                                |
|                   |                                           | 2.2.2.3.1 Develop Pricing Strategy                                   |
|                   |                                           | 2.2.2.3.2 Develop Terms to Manage Risk                               |
|                   |                                           | 2.2.2.3.3 Develop Technical Approach                                 |
|                   |                                           | 2.2.2.3.4 Develop Offer Evaluation Strategy                          |
|                   |                                           | 2.2.2.4 Assess Teaming Options and Partners                          |
|                   |                                           | 2.2.2.4.1 Negotiate Nondisclosure Agreements                         |
|                   |                                           | 2.2.2.4.2 Negotiate Agreements                                       |
|                   |                                           | 2.2.2.4.3 Make Teaming Decisions                                     |
|                   |                                           | 2.2.2.5 Participate in Pre-Offer Conference                          |
|                   |                                           | 2.2.2.6 Finalize Offer                                               |
|                   |                                           | 2.2.2.6.1 Submit Offer and Verify Receipt                            |

**Figure 2. DoDCCM Domain, Competencies, and Tasks for Develop Offer Domain (DoD, 2024)**

Each phase of the contract life cycle is uniquely characterized by a variety of processes that move the contract along the life cycle. The phases build on each other, making it vitally important for the processes within each phase to be conducted as effectively and efficiently as possible. According to the DoDCCM, the pre-award phase consists of all the processes that support contract planning which range from market research and contracting strategy development to requesting bids. The award phase includes all the effort from the buyer and the seller to produce a contract award. Evaluating an offer, selecting a contractor, and debriefing the offerors are a few examples of the activities that occur in this phase. Lastly, the post-award phase consists of all activities and processes that support the administration of the contract until the contract comes to an end and is ready to be closed out. The DoDCCM contends that a contracting workforce that possesses the competencies within this model are key to enabling contracting to deliver the critical needs of the Department.

## Wargaming Overview

The pool of literature that recognizes the usefulness of wargaming both in military (Perla, 1990; Work & Selva, 2015) and non-military (Augier et al., 2018) contexts is extensive. Whether being used as a research (Rubel, 2006), educational (Appleget et al., 2020), learning (U.S. Naval War College, 2026), planning (Lorusso, 2024), strategy development (Augier et al., 2018), training (Perla & McGrady, 2011), or innovation spurring tool (Bartels, 2019; Bartels, 2021), wargaming offers its partakers an opportunity to explore the impact of decisions made in complex and uncertain situations (Augier et al., 2018). As aptly described by Lorusso (2024), “Wargaming provides players with the opportunity to practice their decision making, make critical mistakes, and learn from them, in [a] fail-safe environment, so as, to gain valuable insights when doing so” (p. 6-22). The fortuitous ability to make decisions and learn from the consequences of those decisions in the simulated safe space that a wargame creates,

encourages participants to take risks and be better prepared for real-world conditions (Caffrey, 2019).

Although there is no standard definition of a wargame, the conceptualization of a wargame has evolved and has begun converging on a few common themes (Appleget et al., 2020; U.S. Naval War College, 2026). In 1990, Peter Perla, the author of *The Art of Wargaming* and a recognized wargaming expert (CNA, 2024), offered the following definition of a wargame, “A warfare model or simulation, using rules, data and procedures, not involving actual military forces, and in which the flow of events is affected by, and in turn affects, decisions made during the course of those events by players representing the opposing sides” (p. 274). Over time, Perla updated his definition to better capture the changes in wargaming philosophy since the 1990s (Appleget et al., 2020). His definition, which is similarly reflected in Joint Publication 5-0—Joint Planning Process, is, “A dynamic representation of conflict or competition in a synthetic environment in which people make decision and respond to the consequences of those decisions” (King’s Wargaming Network, 2019; Joint Chiefs of Staff [JCS], 2025). This definition, in conjunction with a variety of others, converge on the conceptualization that “wargames involve people making decisions in a context of competition or conflict (with themselves, other people, or their environment)” (U.S. Naval War College, 2026). Of particular interest in this paper is drawing on the theme of human decision-making in the context of conflict with other people, namely stakeholders, and the environment, namely the external and internal factors that shape an acquisition environment.

Before a wargame can begin to facilitate the examination of human decision-making, the purpose and objectives of the wargame must be clearly articulated (Bartels, 2019). This is a necessary first step to ensure the wargame is designed in an appropriate manner. Appleget et al. (2020), suggests that the main purpose of a game is rooted in either the intent to “convey knowledge or create knowledge” (p. 5) and once that is known, there are three main approaches for attaining that purpose: educational, experiential, and analytical. For example, if the purpose of the game is to have the participants learn from the game, thereby using the game to facilitate imparting knowledge on the participants, then educational and experiential games are the best fit. If the purpose of the game is to extract knowledge from the participants to help generate new information, then analytical games are the most appropriate. It is worth acknowledging that there is a fourth approach called entertainment (Pournelle, 2017); however, that approach is outside of the focus of this paper.

Educational and experiential wargames require the development of learning objectives and linking those learning objectives to the outcomes of the game (Appleget et al., 2020). Through the deliberate design of scenarios, educational wargames have the goal of leaving players with a better understanding of a subject area (Lorusso, 2024) and offer participants the opportunity to apply concepts learned throughout a course, in a realistic but safe environment (Alme & Hvidsten, 2022). Collaboration, creativity, critical thinking, and innovation can also be a byproduct of an educational wargame (Alme & Hvidsten, 2022; Bartels, 2021). Outside of a college or university setting, educational wargames are useful for helping units analyze new areas of operation, evaluate threats from emerging technologies or explore the impacts of new doctrine or proposed force compositions (Appleget et al., 2020).

While similar in purpose to an educational wargame, experiential wargames aim to prepare participants for a specific task or job function (Lorusso, 2024). These games immerse a participant in an environment that replicates the decision-making challenges that the participants may face in their current or future job. A training game is an example of how experiential wargames can be operationalized to help units train their members on standardized job requirements. Experiential games can utilize established training standards as a metric for



measuring the level of proficiency a participant attained from the wargame (Appleget et al., 2020).

Both educational and experiential wargames may be used as precursors to an analytic wargame (Appleget et al., 2020). The insights derived from the participants in an educational or experiential wargame may be intentionally captured, or perhaps unintentionally observed and integrated into a future analytic wargame. Analytic wargames are distinct from educational and experiential wargames because they focus more on the outcome or product that is derived from the decisions made by the wargames' participants, such as contingency plans, courses of action or a concept of operations (Appleget et al., 2020; Lorusso, 2024). An analytic wargame works well for structured problems (Pournelle, 2017) and is characterized as a game that is "designed to collect and analyze information from wargame play, with results feeding directly into a decision, being used to develop additional analytic products or helping to create additional research hypotheses or theories of victory for additional analysis" (Appleget et al., 2020, p. 16). The idea that an analytic game should be an input for further analysis, reemphasizes Perla's (2022) commentary on the cycle of research, to which wargaming is one input of several other activities geared toward better understanding the full scope of a national security challenge.

In addition to wargames having a specific purpose, they are also designed with a certain style in mind. These styles fall on a continuum from seminar style games which are the most flexible and enable the most creativity, to more rigorous games like rigid Kriegsspiel which favor predictability over creativity. Other games that fall along the spectrum are matrix and free Kriegsspiel (Appleget et al., 2020; Pournelle, 2017). Careful thought is placed in determining the appropriate alignment between purpose, style and the wargame's objective(s).

Although the thought of playing a wargame can evoke visions of a high stake, blue verse red military operations game, wargaming's usefulness extends beyond military warfare strategy, to any industry that faces dynamic and complex situations (Fedina & Lucas, 2025). Fedina and Lucas (2025) contend that "from preparing for natural disasters to safeguarding supply chains, wargaming is helping organizations across industries anticipate crises, test strategies, and strengthen their response capabilities" (p.1). This expanded application of wargaming synergizes with the idea that wargaming is inherently focused on human decision-making and can be evaluated within a context where conflict stems from the environment, as could be the case in a natural disaster or humanitarian situation (U.S. Naval War College, 2026). In a similar manner, this paper encourages mission support organizations such as contracting, to utilize the principles of wargaming in a similar manner.

Yet, there is limited academic research that details how mission support functions can incorporate the wealth of wargaming tools into their mission and mission enabling functions. One practical example comes from an Army logistics team who mentions how wargaming was integrated into their logistics operations in AFRICOM (Bennett, 2026). Bennett (2026) describes the concept and need for resilient logistics and discusses the core strategies that enable this concept's proactive foundation. Resilient logistics, defined as "the capacity and capability to absorb, adapt, and respond to disruptions, challenges, and uncertainties in the operational environment, regardless of the cause" (p.2), is proposed as an imperative for helping the logistics function sustain support despite challenges such as, geopolitical changes, dynamic contested environments and unreliable infrastructure. Of the five core strategies that would enable resilient logistics, the strategy of proactive risk management would be when wargaming is implemented. Wargaming would be used to test the contingency plans developed from the risks identified during continuous risk monitoring, with the advantage of testing these plans in advance of an actual crisis situation.



This example underscores how analytical wargaming can be integrated into support function operations. Fictionally extending this example, there could also be educational and experiential games developed to support this shift to resilient logistics. From an educational wargame perspective, a game could be developed around a learning objective that helps logistics professionals understand the challenging and unpredictable nature of the logistics environment for a certain area of responsibility. From an experiential perspective, logistics professionals could engage in a training game with the objective of role playing to understand the challenges of managing joint efforts to maintain logistics operations in a contested logistics environment. Collectively, wargames like these could help support the logistics enterprise's strategic endeavor to shift towards resilient logistics operations.

## Framework

To assist with the development of a wargame, Appleget et al. (2020) propose a five-phase process. The process consists of the following five phases: Initiate, Design, Develop, Conduct, Analyze. The Initiate phase is focused on sponsor engagement and appropriately scoping the wargame to enable the creation of documents such as the Data Collection and Management Plan (DCMP). During the Design phase, the wargame begins to come together with specific details such as the development of the scenarios and adjudication methods. In the third stage, Develop, the wargame's initial design is tested, and through an iterative process with elements of the design phase, the wargame is ready for execution. During the Conduct phase, the wargame is played and once complete, a post wargame analysis and finalization of the findings round out the Analysis phase. For specific details on the key tasks for each phase reference, *The Craft of Wargaming* by Appleget et al. (2020).

While initially developed for analytic wargames, this five-phase process can be used for the development of educational and experiential wargames. Appleget et al. (2020) offer guidance on how to adapt the five phases for each game. Table 1 shows both overlapping and unique tasks for each phase where the underlined tasks are those specific to either an educational or experiential wargame. The preponderance of differences between analytic and educational/experiential wargames occurs during the initiate phase. Instead of having a sponsor to guide the objective(s) of the wargame, learning objectives for educational games and tasks for experiential games serve as the foundation of the game. During the design phase, pre- and post-wargame assessments should be designed to measure the extent to which the wargame enabled the participants to achieve educational or experiential objectives. In the conduct phase, a pre-wargame assessment helps to establish the baseline knowledge of the participants while a debrief at the conclusion of the game is an opportunity to revisit any learning objectives or tasks that were not sufficiently addressed. Lastly, the pre-wargame and post-wargame assessment are evaluated during the analysis phase.



**Table1. Five Phases of Wargame Development**  
(adapted from Appleget et al., 2020)

| Phase    | Educational                                                                                                                                                                                                                                                                                                                        | Experiential                                                                                                                                                                                                                                                                                                                                                             |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Initiate | <ul style="list-style-type: none"> <li>• <u>Goal: Build wargame around learning objectives</u></li> <li>• Form core wargaming team as applicable</li> <li>• Scope problem</li> <li>• <u>Develop learning objectives</u></li> <li>• <u>Develop DCMP linked to learning objectives</u></li> </ul>                                    | <ul style="list-style-type: none"> <li>• <u>Goal: Build wargame around the intended participant experience</u></li> <li>• <u>Problem scoping includes identifying tasks (training tasks if doing a training wargame) that the wargame will center around</u></li> <li>• Form core wargaming team as applicable</li> <li>• <u>Develop DCMP linked to tasks</u></li> </ul> |
| Design   | <ul style="list-style-type: none"> <li>• Determine scenario(s)</li> <li>• Select adjudication approach (models, methods, tools)</li> <li>• Determine player roles</li> <li>• Determine wargame data requirements</li> <li>• Iterate with development phase as needed</li> <li>• <u>Pre- and post-wargame assessment</u></li> </ul> | <ul style="list-style-type: none"> <li>• Determine scenario(s)</li> <li>• Select adjudication approach (models, methods, tools)</li> <li>• Determine player roles</li> <li>• Determine wargame data requirements</li> <li>• Iterate with development phase as needed</li> <li>• <u>Pre- and post-wargame assessment</u></li> </ul>                                       |
| Develop  | <ul style="list-style-type: none"> <li>• Play-test all components of wargame (multiple times as necessary)</li> <li>• Blind play-test wargame</li> <li>• Full dress rehearsal</li> </ul>                                                                                                                                           | <ul style="list-style-type: none"> <li>• Play-test all components of wargame (multiple times as necessary)</li> <li>• Blind play-test wargame</li> <li>• Full dress rehearsal</li> </ul>                                                                                                                                                                                 |
| Conduct  | <ul style="list-style-type: none"> <li>• <u>Pre-wargame assessment</u></li> <li>• Collect data</li> <li>• Manage players</li> <li>• Exercise contingencies (as necessary)</li> <li>• <u>Debrief</u></li> </ul>                                                                                                                     | <ul style="list-style-type: none"> <li>• <u>Pre-wargame assessment</u></li> <li>• Collect data</li> <li>• Manage players</li> <li>• Exercise contingencies (as necessary)</li> <li>• <u>Debrief</u></li> </ul>                                                                                                                                                           |
| Analyze  | <ul style="list-style-type: none"> <li>• <u>Post-wargame assessment</u></li> </ul>                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• <u>Post-wargame assessment</u></li> </ul>                                                                                                                                                                                                                                                                                       |

The five-phase wargame development process in Table 1, provides a solid framework for overlaying contracting specific educational and experiential wargame objectives. To this end, Table 2 provides examples of pre-award, award, and post-award wargame objectives. These objectives are separated into example objectives for educational, experiential, and analytical wargames. The educational and training objectives stem from the activities and competencies associated with each phase of the contract life cycle and as outlined in the DoDCCM (2024). For example, Task 2.1.1.3 Perform Risk Assessment in the Pre-Award phase requires contracting professionals to perform a risk analysis related to the development of a solicitation. To achieve that task, an experiential wargame can be developed that immerses a contingency contracting officer into a contingency environment where they are provided scenarios centered around conducting risk analysis activities. Each example objective also has a target audience to help frame the level of the game in terms of tactical, operational or strategic focus.



**Table 2. Wargame Objectives Across the Contract Life Cycle**

| Contract Life Cycle Stage | Wargame Type | Example Objective                                                                                                                                                                          | Target Audience                                |
|---------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|
| Pre-Award                 | Educational  | Apply business acumen to evaluate the impact of market trends, geopolitical issues and Defense objectives to battle test an acquisition plan                                               | Junior to mid-level contracting professionals  |
|                           | Experiential | Demonstrate the ability to conduct a risk analysis in a contingency environment                                                                                                            | Contingency contracting officers               |
|                           | Analytical   | Determine which theater support contract approach, is most effective for a new Area of Operations (AO)                                                                                     | Senior leaders                                 |
| Award                     | Educational  | Demonstrate an understanding of how a contract award affects a mission partners ability to carry out the mission in a contingency environment                                              | Contingency contracting officers               |
|                           | Experiential | Role play as a contingency contracting officer and identify indicators of fraud and malicious intent during source selection evaluation                                                    | Contingency contracting officers               |
|                           | Analytical   | Analyze the effectiveness of a novel contract award policy                                                                                                                                 | Senior leaders                                 |
| Post-Award                | Educational  | Demonstrate understanding of the impact of cost, schedule and performance tradeoffs on stakeholders for a given contract action                                                            | Junior to mid-level contracting professionals  |
|                           | Experiential | Demonstrate the ability to identify at least two courses of action for a post award crisis that make good business sense and calculate/mitigate the risks of the various courses of action | Junior to mid-level contracting professionals  |
|                           | Analytical   | In collaboration with government partners, test contingency plans for alternate sourcing strategies in a contested logistics environment                                                   | Mid- to senior level contracting professionals |

### Example Wargame

As previously mentioned, there are numerous activities that occur throughout the contracting life cycle. One such activity that occurs during the pre-award phase is the development of the acquisition plan. An acquisition plan is the road map for an acquisition and



“must address all the technical, business, management, and other significant considerations that will control the acquisition” (Federal Acquisition Regulation 7.105, 2026). An educational wargame centered around stress testing the acquisition plan before it is employed is the main objective of the game. In this example, the opposing forces against the acquisition plan stem from the acquisition environment and include factors such as, but not limited to market trends, geopolitical issues, intellectual property challenges, and supply chain risks. Extensive research into each of these factors is conducted to properly contextualize the scenario.

Utilizing the five-phase framework for an educational wargame, Table 3 provides an overview of the development of the following scenario: The Air Force FLYT office is procuring communications equipment and associated installation/integration services for the STARS Communications Program. A draft acquisition strategy is provided to each team. In the Initiate phase, two lesson objectives (LO) are identified that are focused on testing the acquisition plan and updating the plan as gaps are identified. Data collection occurs in the form of a decision document that each team populates with the rationale for their decisions along with any assumptions that were made. The final deliverable is an updated acquisition plan. During the Design phase, the scenario and vignettes are designed to leverage extensive research conducted on market trends, geopolitical issues, and Defense objectives with the potential to impact communications equipment and services. The facilitator is the adjudication method, and both the pre-award and post-award assessments are designed. In the Development phase, play-testing, blind-testing, and a full rehearsal of all elements of the game is conducted. During the Conduct phase, pre-wargame assessments were collected the day before, and the game commences. At the end of the Conduct phase, a debrief is provided by the students summarizing the changes that were made to the acquisition plan followed by an open discussion facilitated by the faculty member. The Analysis phase concluded the game with a post-wargame assessment to assist with analyzing how much learning was achieved. It is worth noting that the post-wargame assessment could also be completed before the discussion with the faculty member.

**Table 3. Example Wargame Development**

| Phase    | Educational                                                                                                                                                                                                                                                                                     | Notional Application                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Initiate | <ul style="list-style-type: none"> <li>• <u>Goal: Build wargame around learning objectives</u></li> <li>• Form core wargaming team as applicable</li> <li>• Scope problem</li> <li>• <u>Develop learning objectives</u></li> <li>• <u>Develop DCMP linked to learning objectives</u></li> </ul> | <ul style="list-style-type: none"> <li>• <u>Goal: Build wargame around learning objectives</u></li> <li>• Core team consists of faculty member and assistant</li> <li>• Acquisition plans are an important part of the acquisition process, yet limited time is spent testing these strategies to see to what extent they can handle a variety of potential future issues</li> <li>• LO 1: Apply business acumen to evaluate the impact of market trends, geopolitical issues and Defense objectives to battle test an acquisition plan</li> <li>• LO 2: Identify gaps in an acquisition plan and update the plan with appropriate information to fill the gaps</li> <li>• Data collection includes the use of a decision document provided to each</li> </ul> |



|         |                                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|         |                                                                                                                                                                                                                                                                                                                        | team to capture the rationale for each decision, noting any assumptions. Collection also stems from the final deliverable of the revised acquisition plan.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Design  | <ul style="list-style-type: none"> <li>Determine scenario(s)</li> <li>Select adjudication approach (models, methods, tools)</li> <li>Determine player roles</li> <li>Determine wargame data requirements</li> <li>Iterate with development phase as needed</li> <li><u>Pre- and post-wargame assessment</u></li> </ul> | <ul style="list-style-type: none"> <li>Main Scenario: The Air Force FLYT office is procuring communications equipment and associated installation/integration services for the STARS Communications Program. A draft acquisition strategy is provided to each team.</li> <li>Vignettes: Developed from extensive research into market trends, geopolitical issues and Defense objectives with the potential to impact communications equipment and services</li> <li>Facilitator adjudication method</li> <li>Contracting Officer for procurement of the STARS requirement. Participants are placed in teams to leverage varying degrees of experience</li> <li>Participants provided with an acquisition strategy and computer access to the FAR and other pertinent regulations</li> <li>Iterate with development phase as needed</li> <li>Pre-award assessment provides insight into each player's level of experience with developing acquisition plans and experience in contracting in general</li> <li>Post-award assessment provides insights into how much each participant learned from the game through post-game interviews and debriefings</li> </ul> |
| Develop | <ul style="list-style-type: none"> <li>Play-test all components of wargame (multiple times as necessary)</li> <li>Blind play-test wargame</li> <li>Full dress rehearsal</li> </ul>                                                                                                                                     | <ul style="list-style-type: none"> <li>Play-test vignettes, game flow, and contingency plans</li> <li>Conduct test with non-participants with contracting experience</li> <li>Full dress rehearsal</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Conduct | <ul style="list-style-type: none"> <li><u>Pre-wargame assessment</u></li> </ul>                                                                                                                                                                                                                                        | <ul style="list-style-type: none"> <li>Provided participants with the pre-wargame assessment the day prior to the game to allow time to process the information</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



|         |                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|         | <ul style="list-style-type: none"> <li>• Collect data</li> <li>• Manage players</li> <li>• Exercise contingencies (as necessary)</li> <li>• <u>Debrief</u></li> </ul> | <ul style="list-style-type: none"> <li>• Collect data throughout the game</li> <li>• Once facilitator per team</li> <li>• Exercise contingencies (as necessary)</li> <li>• Participants debriefed a summary of the changes to their strategy along with their rationale. Upon completion of the debriefs the faculty member provides an overview of the game's objectives, facilitates an open discussion of the overall lessons learned</li> </ul> |
| Analyze | <ul style="list-style-type: none"> <li>• <u>Post-wargame assessment</u></li> </ul>                                                                                    | <ul style="list-style-type: none"> <li>• Participants given post-wargame assessments and analysis between the pre- and post-wargames assessments is analyzed</li> </ul>                                                                                                                                                                                                                                                                             |

## Future Research

There are several opportunities for future research that can continue the discussion on how wargaming can be utilized in not only contracting but more broadly, acquisitions. The acquisition team consists of multiple players across functional specialties (FAR 1.102-4, 2026). Wargames that include members of the entire acquisition team could facilitate collective decision-making and collaboration. Future research could also examine how acquisition wargames could inform an acquisition focused exercise thereby creating what Peter Perlas called a cycle of research (Perla, 2022). Furthermore, exploring how to structure a wargaming ecosystem as part of the training and equipping of contracting and acquisitions professionals is a worthwhile endeavor. Additionally, another area for future research is to examine opportunities to create wargames that may not fall directly within one phase of the contracting life cycle. While the contracting life cycle model is useful for linking tasks to objectives, a sponsor may be interested in understanding organizational or manpower problems which may not be directly captured in the contracting life cycle

As with any tool, wargaming has inherent limitations. A wargame is most useful when the question or objective is exploratory and not seeking validation (Bartels, 2019). Ensuring there is alignment between the desired objectives, and the game design is an imperative which may require a significant time investment to complete the appropriate level of iteration and testing (Applegate, 2020). Nonetheless, the advantage of creating a fail-safe environment to test human-decision making is worth the investment.

## Conclusion

"Innovation thrives in a culture that embraces experimentation and tolerates—better yet, encourages—dissent and risk-taking. We must create an environment in the Department of Defense that encourages exactly this type of thinking. Building a reinvigorated wargaming enterprise is a major step toward that goal" (Work & Selva, 2015, p. 3).

This quote embodies one of the intents of this paper, to examine how wargaming could be incorporated into the contracting enterprise as a tool for spurring innovation while creating a safe place for risk-taking. Given the type of competencies required of contracting professionals, developing educational, experiential and/or analytic wargames may be a useful investment in the training and equipping ecosystem for contracting professionals. The proposed framework shows how wargames can be crafted to meet certain objectives or accomplish required training tasks. Overall, the complex and dynamic Defense acquisition landscape is driving the need for



contracting professionals to exercise the competencies in the DoDCCM. Integrating wargaming into the contracting enterprise can stimulate critical thinking and prepare contracting professionals for future challenges.

## References

- Alme, V., & Hvidsten, A. (2022). To learn or not to learn: On the importance of mode switching in educational wargames. *Scandinavian Journal of Military Studies*, 5(1), 269–281. <https://doi.org/10.31374/sjms.123>
- Appleget, J., & Burks, R. (2020). *The craft of wargaming*. Naval Institute Press.
- Assistant Secretary of War (Acquisition). (n.d.). *Contract Policy*. Defense Pricing, Contracting, and Acquisition Policy (DPCAP) <https://www.acq.osd.mil/asda/dpc/cp/policy/workforce-development.html>
- Augier, M., Dew, N., Knudsen, T., & Stieglitz, N. (2018). Organizational persistence in the use of war gaming and scenario planning. *Long Range Planning*, 51(4), 511–525. <https://doi.org/10.1016/j.lrp.2017.12.005>
- Bartels, E. (2019). *Getting the most out of your wargame: Practical advice for decision-makers*. War on the Rocks. <https://warontherocks.com/2019/11/getting-the-most-out-of-your-wargame-practical-advice-for-decision-makers/>
- Bartels, E. M. (2021). *Wargames as an educational tool*. Rand. <https://www.rand.org/pubs/commentary/2021/02/wargames-as-an-educational-tool.html>
- Bennett, M. (2026). *Resilient logistics: A blueprint for success in contested environments*. [https://www.army.mil/article/290281/resilient\\_logistics\\_a\\_blueprint\\_for\\_success\\_in\\_contested\\_environments](https://www.army.mil/article/290281/resilient_logistics_a_blueprint_for_success_in_contested_environments)
- Caffrey, M. B., Naval War College (U.S.). Press, & Naval War College (U.S.). Center For Naval Warfare Studies. (2019). *On wargaming: How wargames have shaped history and how they may shape the future*. Naval War College Press.
- Commerce and Contract Management Institute. (2019). *Contract management standard*. [https://cdn.prod.website-files.com/682219f2368e74799ec8747e/6852f64457c859a0dceab1e2\\_CCM-Institute-CMS-Edition4.pdf](https://cdn.prod.website-files.com/682219f2368e74799ec8747e/6852f64457c859a0dceab1e2_CCM-Institute-CMS-Edition4.pdf)
- CNA. (2024). *CNA wargaming reports by Peter Perla*. <https://www.cna.org/centers-and-divisions/cna/ow/gaming-and-integration/peter-perla>
- DoD. (2024). *Department of Defense Contracting Competency Model (DoDCCM)*. [https://www.acq.osd.mil/asda/dpc/cp/policy/docs/wd/DoD-Contracting-Competency-Model-\(2024\).pdf](https://www.acq.osd.mil/asda/dpc/cp/policy/docs/wd/DoD-Contracting-Competency-Model-(2024).pdf)
- FAR 1.102-4, Acquisition Team (2026). [https://www.acquisition.gov/far/part-1#FAR\\_1\\_602\\_1](https://www.acquisition.gov/far/part-1#FAR_1_602_1)
- FAR 1.602, Contracting Officers (2026). [https://www.acquisition.gov/far/part-1#FAR\\_1\\_602\\_1](https://www.acquisition.gov/far/part-1#FAR_1_602_1)
- FAR 7.105, Contents of written acquisition plans (2026). [https://www.acquisition.gov/far/part-7#FAR\\_Subpart\\_7\\_1](https://www.acquisition.gov/far/part-7#FAR_Subpart_7_1)
- Fedina , K., & Lucas, R. (2025). *Building societal resilience through wargaming*. RAND Corporation. <https://www.rand.org/pubs/commentary/2025/04/building-societal-resilience-through-wargaming.html>





- King's Wargaming Network. (2019, March 29). *Peter Perla, "The art & science of wargaming to innovate & educate"* [Video]. YouTube. <https://www.youtube.com/watch?v=rxLQmPA1-4o>
- Lorusso, L. (2024). *Introduction to wargaming*. NATO. <https://publications.sto.nato.int/publications/STO%20Educational%20Notes/STO-EN-SAS-195/EN-SAS-195-06.pdf>
- Perla, P. (1990). *The art of wargaming : A guide for professionals and hobbyists*. Naval Institute Press.
- Perla, P. (2022). Wargaming and the cycle of research and learning. *Scandinavian Journal of Military Studies*, 5(1), 197–208. <https://doi.org/10.31374/sjms.124>
- Perla, P., & McGrady, E. (2011). Why wargaming works. *Naval War College Review*, 64(3). <https://digital-commons.usnwc.edu/nwc-review/vol64/iss3/8>
- Pournelle, P. (2017). Designing wargames for the analytic purpose. *Phalanx*, 50(2), 48–53. <https://www-jstor-org.nps.idm.oclc.org/stable/pdf/26296384>
- Rubel, R. C. (2006). The epistemology of war gaming. *Naval War College Review*, 59(2), Article 8.
- U.S. Naval War College. (2026). *About wargaming*. <https://www.usnwc.edu/Research-and-Wargaming/Wargaming/About-Wargaming>
- Work, B., & Selva, P. (2015). *Revitalizing wargaming is necessary to be prepared for future wars*. War on the Rocks. <https://warontherocks.com/2015/12/revitalizing-wargaming-is-necessary-to-be-prepared-for-future-wars/>



# Is the Military Attracting New Companies? Assessing First-Time Entrants into the Defense Market

**Amanda Bresler**—serves as President of PW Communications, Inc. She runs SHELDON, a subsidiary of PW Communications that provides custom analytics products and data-driven advisory services to federal and commercial clients. Prior to joining PW Communications, she worked as Chief Operating Officer for Maurice Cooper Brands. She serves on the board of directors of PW Communications; St. Dalfour SAS, a French food company; and Chatham International Inc. She graduated cum laude from Georgetown University's McDonough School of Business. [abresler@pwcommunications.com]

**Alex Bresler**—serves as Chief Data Officer for PW Communications Inc. He is a data-driven technologist, investor, and advisor to early stage venture funded technology companies. He is a data scientist and programmer with experience supporting clients in defense, financial services, law, real estate, and sports. He graduated from the Wharton School at the University of Pennsylvania. [alexbresler@pwcommunications.com]

**David Lipa**—is a software engineer and advisor for PW Communications, Inc. He is the founder of multiple successful companies, an angel investor in the aerospace industry, and an advisor to technology startups. He has broad experience with venture finance, valuation, diligence standards and software engineering. David graduated from Stanford University with a BSc in Management Science & Engineering. [dlipa@pwcommunications.com]

## Abstract

Over the last 20+ years, companies outside of the traditional defense industrial base (DIB) increasingly drive innovation in areas critical to national security. This paradigm shift has resulted in the Department of War (DoW) investing billions of dollars into innovation programs, rapid acquisition programs, small business outreach programs, and accelerators (“innovation initiatives” or “initiatives”) that have a stated purpose of making it easier for innovative commercial companies to break into the defense market. In our 2020 Naval Postgraduate School paper, we assessed contract award data from fiscal year (FY) 2010 through FY2019 and determined that, rather than expanding the industrial base, these initiatives primarily benefit entrenched defense contractors. In this paper, we revisited the analysis of new versus existing DoW vendors, using updated datasets from the Federal Procurement Data System (FPDS) between FY2015 and FY2024. We determined that the DoW remains largely inaccessible for nonentrenched contractors, and the majority of first time vendors to the defense market do not purvey commercial technologies. In addition to quantitatively analyzing the types of companies that sell to the DoW, we discuss the underlying issues that have made the defense market so inhospitable towards new entrants. Finally, we offer a series of recommendations for how the DoW can better attract, engage, and retain innovative commercial companies.

## Introduction

The Department of War (DoW) faces a dual imperative in managing the defense industrial base (DIB): attracting and retaining suppliers best positioned to meet warfighter needs, while consolidating where strategic efficiency demands it. Dozens of DoW innovation programs, rapid acquisitions programs, small business outreach programs, and accelerators (“innovation programs”) receive billions in funding with the stated objective of accelerating the adoption of innovative new technologies forcewide. Simultaneously, the DoW has shifted to consolidated purchasing strategies for categories like parts supply, field maintenance activities, and performance-based logistics to improve readiness (AAR Corp).

However, the DoW has not maintained a common operating picture of its supplier base. As a result, there has been a lack of useful metrics to track changes in the DIB in general, and to determine the effectiveness of these programs in particular. Our research aimed to fill this gap in two primary ways. First, by using public data to analyze trends in the DIB, the types of



companies entering and exiting the DIB over the last decade, and how new entrants break into the defense market. Secondly, by demonstrating, via our research process, the feasibility of using public data and readily-available tools to build and maintain a detailed view of the DIB. In addition to the policy-focused recommendations we offer, this paper serves as a technical framework for the DoW to implement data-driven processes in its overall management and evaluation of the industrial base.

## Sizing the Defense Industrial Base

Before analyzing changes to the composition of the DIB over time, we needed to isolate DoW vendor data by year. To do so, we aggregated contract award data from the Federal Procurement Data System (FPDS), the centralized, real-time database for government procurement transactions. We then filtered the data to isolate contracts funded and awarded by the DoW from fiscal year (FY) 2015 through FY2024.<sup>1</sup> To determine the unique number of companies associated with these contract actions, we grouped each contract action by its associated Unique Entity Identifier (UEI). The resulting 142,890 UEIs became our portfolio for analysis. We then leveraged large language models (LLMs) and generative artificial intelligence (GenAI) to build user-friendly dashboards to analyze and visualize this data.

## The Declining DIB

As shown in Figures 1, 2, and 3 the DIB contracted nearly 30% from FY2015 to FY2024, from 67,797 unique vendors to 48,035. Simultaneously, defense spending surged from \$275 billion in FY2015 to \$462 billion in FY2024.

Total Vendor Pool & Defense Obligations, FY2015–FY2024

| FY     | Actions | Vendors | Total Obligated |
|--------|---------|---------|-----------------|
| FY2015 | 3.2M    | 67,797  | \$274.7B        |
| FY2016 | 3.6M    | 65,340  | \$295.7B        |
| FY2017 | 3.7M    | 63,537  | \$323.7B        |
| FY2018 | 4.4M    | 60,419  | \$357.9B        |
| FY2019 | 4.4M    | 57,374  | \$393.9B        |
| FY2020 | 4M      | 53,965  | \$436.2B        |
| FY2021 | 4.3M    | 52,281  | \$400.0B        |
| FY2022 | 4.3M    | 48,700  | \$422.4B        |
| FY2023 | 4.4M    | 48,554  | \$472.3B        |
| FY2024 | 4.4M    | 48,035  | \$461.7B        |

Source: FPDS, defense-funded and awarded. Actions = FPDS transaction rows. Obligations are net.

**Figure 1. Sizing the Defense Market Annually**

<sup>1</sup> For the purposes of our analysis and to remain consistent, we focus on contract actions funded and awarded by the DoW. There are instances in which DoW may serve as the awarding entity, but funding has come from another department; and/or instances in which the DoW funds something that is awarded by a non-DoW agency. Those instances are not included in our analysis.





**Figure 2. Unique DoW Vendors by Year**

Nominal and inflation-adjusted (2015 base year) obligations

| Fiscal Year | Vendors | Contracts | Nominal  | Real (2015) |
|-------------|---------|-----------|----------|-------------|
| FY2015      | 67,797  | 659,617   | \$274.7B | \$274.7B    |
| FY2016      | 65,340  | 722,077   | \$295.7B | \$292.0B    |
| FY2017      | 63,537  | 781,094   | \$323.7B | \$313.0B    |
| FY2018      | 60,419  | 3,976,619 | \$357.9B | \$337.9B    |
| FY2019      | 57,374  | 4,013,880 | \$393.9B | \$365.2B    |
| FY2020      | 53,965  | 3,637,604 | \$436.2B | \$399.4B    |
| FY2021      | 52,281  | 3,968,365 | \$400.0B | \$349.9B    |
| FY2022      | 48,700  | 3,961,646 | \$422.4B | \$342.1B    |
| FY2023      | 48,554  | 4,044,108 | \$472.3B | \$367.4B    |
| FY2024      | 48,035  | 4,048,205 | \$461.7B | \$349.3B    |

Source: FPDS, defense-funded and awarded. CPI-U deflator (BLS).

**Figure 3. Defense Spending, FY2015–FY2024**

### Research Limitation: Inflated Vendor Count

Large corporations frequently register multiple UELs, each of which our methodology treats as a discrete vendor, in turn inflating our overall count of suppliers. In the new entrants section, we employ a process to flag and consolidate affiliated UELs under a single parent entity. We recommend applying this approach to the full dataset in future research.

### Growth for a Few at the Expense of Many

Rather than undermine the results, this research limitation heightens the directional trends revealed in the data. Based on the significant growth in defense spending and concurrent



decline in vendors, a subset of defense contractors have materially expanded their market share over the last decade. These “winners” reap the financial rewards of a consolidating DIB, risk material revenue loss if expansion and diversification efforts succeed, and have outsized resources and influence to direct towards setting and maintaining industrial base policy norms. The majority of these winners are private companies and thus motivated primarily by revenue generation. Some have investment from the public or private capital markets, which means they have an explicit obligation to maximize shareholder value and/or increase profitability. These incentives can be at odds with national security imperatives, particularly in critical or vulnerable areas of the supply chain. These dynamics must inform any detailed analysis of DIB health, including the one that follows.

### **Composition of the Defense Industrial Base**

DIB policy must reflect a balance between the need for strategic consolidation in certain categories, and the need to retain existing vendors and/or introduce new vendors in others. To assess the types of companies that entered and exited the defense market during our analysis period, we began by stratifying the DIB data in each year, across three metrics:

- 1) Count of existing vendors (companies active in the defense market, prior to FY2015)
- 2) Count of new vendors (companies with no prior DoW funding)
- 3) Count of vendors that departed the defense market (companies for which a given fiscal year was its last year as an active supplier to DoW)

From there, we could closely analyze the types of companies in each category.

### **Stratifying the DIB: Technical Approach**

For each unique DoW vendor in our analysis group, we isolated the first recorded defense funded-and-awarded contract action (“DoW funding”), dating back to 1977.<sup>2</sup> We treated entities with defense funding prior to FY2015 as “existing DoW” vendors. Those with no DoW funding prior to FY2015 were treated as “new” in the fiscal year of their first DoW funded-and-awarded contract action. To calculate “departing” vendors, we looked at the most recent DoW payment each company received. We categorized a company as a “departed vendor” if it received DoW funding at some point during our analysis period, but not in subsequent fiscal years.

### **Research Limitations & Future Research**

For the purposes of this research, we leveraged prime contract award data from FPDS. We recommend future research incorporate subcontract award data from USASpending to capture a more complete picture of the DIB. Additionally, we recognize that a vendor may cease working with the DoW at a point in time and reengage at a future date. Constraints to the timeframes in our analysis may inflate “departing vendor” data, particularly for later fiscal years where vendors may have had only one year of dormancy. The DoW must track these metrics over time, continuously adjusting entry and exit numbers as new data becomes available.

### **Composition of the DIB: FY2015-FY2024**

From FY2015-FY2024, a total of 59,189 vendors entered the DIB for the first time, while 94,849 vendors departed. For every new company that entered the defense market, approximately two vendors left. If this trend continues, DoW stands to lose another 10,000+ net vendors by FY2030.

---

<sup>2</sup> FPDS data begins in 1977





Figure 4 shows a breakdown of new, existing, and departing vendors in each year. The “Count of Departing DoW Vendors” column refers to the total number of vendors for whom that fiscal year was their last recorded defense funded-and-awarded contract action during our analysis period. For instance, 11,818 vendors had defense funded-and-awarded contract actions in FY2015, but not FY2016-FY2024. We exclude FY2024 from the departure analysis because we do not incorporate FY2025 data, and therefore cannot confirm how many vendors returned or exited in FY2025.

Total, existing, new, and departing vendors by fiscal year

| Fiscal Year | Total  | Existing | New to DoW | Departing |
|-------------|--------|----------|------------|-----------|
| FY2015      | 67,797 | 59,911   | 7,886      | 11,818    |
| FY2016      | 65,340 | 57,914   | 7,426      | 11,340    |
| FY2017      | 63,537 | 56,534   | 7,003      | 11,359    |
| FY2018      | 60,419 | 54,052   | 6,367      | 10,835    |
| FY2019      | 57,374 | 51,609   | 5,765      | 10,438    |
| FY2020      | 53,965 | 48,267   | 5,698      | 9,432     |
| FY2021      | 52,281 | 47,135   | 5,146      | 9,709     |
| FY2022      | 48,700 | 44,606   | 4,094      | 8,682     |
| FY2023      | 48,554 | 43,833   | 4,721      | 11,236    |
| FY2024      | 48,035 | 42,955   | 5,080      | —         |

New to DoW = first defense-funded action in that FY. FY2024 departing blank by design.

**Figure 4. Stratifying the DIB**

### Who Left? Oldtimers vs. Newcomers

Stratifying the DIB in this way provided meaningful summary statistics about the defense market in general, and also allowed for a more nuanced analysis of DIB participants by type. We started by closely analyzing the features of departing companies, beginning with their tenure in the defense market. For each departing company, we joined its earliest DoW funding action from FPDS, which dates back to 1977, and compared it to the date of its last DoW award. We then split departed vendors into two groups: 1) “single-year” vendors, active in the defense market for only one year; and 2) “multi-year” vendors that had sustained activity in the defense market before departing.

As shown in Figure 5, approximately 20% of departing companies during our analysis period were single-year vendors, whereas 80% of departing companies were multi-year firms. This breakdown highlights two points: a sizable share of companies enter the defense market but do not stay, which will serve as an important consideration during our analysis of new entrants; and the majority of departing firms are legacy providers to the DoW.



Figure 5: Exiting Firms — Single-Year vs Multi-Year

Vendors whose last defense action was in the given fiscal year

| Last Active FY | Total  | Single - Year | Multi - Year |
|----------------|--------|---------------|--------------|
| FY2015         | 11,818 | 2,638 (22%)   | 9,180 (78%)  |
| FY2016         | 11,340 | 2,449 (22%)   | 8,891 (78%)  |
| FY2017         | 11,359 | 2,445 (22%)   | 8,914 (78%)  |
| FY2018         | 10,835 | 2,280 (21%)   | 8,555 (79%)  |
| FY2019         | 10,438 | 1,798 (17%)   | 8,640 (83%)  |
| FY2020         | 9,432  | 1,879 (20%)   | 7,553 (80%)  |
| FY2021         | 9,709  | 1,748 (18%)   | 7,961 (82%)  |
| FY2022         | 8,682  | 1,409 (16%)   | 7,273 (84%)  |
| FY2023         | 11,236 | 1,980 (18%)   | 9,256 (82%)  |

*Figure 5. Exiting Firms—Single-Year vs. Multi-Year*

### What Did Departing Firms Sell?

Particularly in the context of these long-tenured vendors, the implications of attrition depend heavily on what goods/services departing firms had provided. Simply put, losing pencil manufacturers is of different consequence than losing manufacturers of specialized aerospace components.

To better understand the types of companies that departed, we explored the entirety of the contract data associated with departed vendors dating back to 1977. Specifically, we looked at companies' associated North American Industry Classification System (NAICS) codes and Product Service Codes (PSCs). A NAICS code is a label used by the government to classify the type of work a firm does, and a PSC is a code used by the government to describe the product or service being purchased via a specific contract. Of note, a single departed vendor can be associated with multiple PSCs and NAICS.

Figure 6 shows the top 15 industries associated with departing vendors, based on NAICS code. Figures 7 and 8 detail the top 15 PSCs associated with departing vendors. The majority of departing vendors by PSC were hotels, which also favored heavily in the NAICS data. Attrition in lodging is unlikely to portend supply chain vulnerability. However, the largest share of departing vendors by NAICS—about 6% of all companies that exited—were associated with NAICS 541712 and 541715, both of which cover “Research and Development (R&D) in the Physical, Engineering, and Life Sciences (Except Nanotechnology and Biotechnology)” (U.S. Census Bureau). Like all sectors, some churn among R&D firms should be expected. However, the DoW must monitor this supplier erosion carefully to understand what causes these firms to leave. Do their technologies fail to meet mission needs, or are they casualties of the government’s burdensome bureaucracy?






**Figure 6. Top 15 NAICS Industries by Departed Vendor Count**




**Figure 7. Top 15 Products/Services Represented by Departing Firms, by PSC**



Most common categories among vendors who left the defense market

| PSC  | Description                                | Vendors | Obligated  |
|------|--------------------------------------------|---------|------------|
| V231 | LODGING - HOTEL/MOTEL                      | 3,707   | \$702.1M   |
| R499 | OTHER PROFESSIONAL SERVICES                | 3,018   | \$2,675.8M |
| 6515 | MEDICAL AND SURGICAL INSTRUMENTS, EQUIPMEN | 2,656   | \$2,971.3M |
| 7030 | ADP SOFTWARE                               | 2,623   | \$856.9M   |
| 5340 | HARDWARE                                   | 2,343   | \$191.9M   |
| R425 | ENGINEERING AND TECHNICAL SERVICES         | 1,783   | \$3,462.4M |
| 6640 | LABORATORY EQUIPMENT AND SUPPLIES          | 1,702   | \$776.1M   |
| AD92 | OTHER DEFENSE - APPLIED RESEARCH/EXPLORATO | 1,416   | \$1,120.6M |
| 7110 | OFFICE FURNITURE                           | 1,387   | \$339.2M   |
| U099 | OTHER EDUCATION AND TRAINING SERVICES      | 1,374   | \$418.3M   |
| 5999 | MISCELLANEOUS ELECTRICAL AND ELECTRONIC CO | 1,337   | \$155.8M   |
| 5895 | MISCELLANEOUS COMMUNICATION EQUIPMENT      | 1,311   | \$362.5M   |
| J099 | MAINTENANCE, REPAIR AND REBUILDING OF EQUI | 1,244   | \$253.3M   |
| AD25 | SERVICES - OPERATIONAL SYSTEMS DEVELOPMENT | 1,155   | \$10.0M    |
| Z2JZ | REPAIR OR ALTERATION OF MISCELLANEOUS BUIL | 1,133   | \$2,652.5M |

**Figure 8. Top 15 PSCs—Departed Vendors**

Looking at the PSC data, losses in categories like Professional Services, Medical Instruments/Equipment, Automatic Data Processing (ADP) Software, and Hardware, could also be consequential. In Figure 8, the column “Total Obligations” provides the total defense funding captured by departing vendors in a given PSC between FY2015-FY2024. For example, the 2,656 suppliers of Medical/Surgical Equipment that departed the DIB between FY2015-FY2024 received more than \$2.9 billion in defense funding during that time.

For high-value losses such as Medical/Surgical Equipment and ADP Software, the DoW must understand what suppliers replaced those departed vendors; if and how the outflow of suppliers strained the remaining industrial base; and why those suppliers left. Attention must also be paid to the 2,343 “Hardware” vendors that departed—purveyors of literal nuts and bolts—that comprised only \$192 million in obligations. In other words, thousands of suppliers accounting for relatively low contract values ceased operating in the defense market. To what extent did they provide speciality parts? Do they remain active in the commercial market, or have they shuttered entirely? Were they driven out of the DoW as a result of contract bundling by the Defense Logistics Agency (DLA)? If so, did forced consolidation produce material gains in efficiency or cost for DoW?

### **Context is Key: Qualitative Research Needs**

These questions highlight an important analytical limitation: the data cannot tell us why a company exits the DIB. Potential factors range from uncertainty in defense spending, ballooning compliance costs, and the comparative ease of working with commercial customers to considerations entirely unrelated to the dynamics of the defense market (Serbu). Closing this knowledge gap will require qualitative insights that no dataset alone can provide. However,



publicly available data, LLMs, and GenAI afford new ability to manage the quantitative aspects of the DIB. The DoW should leverage these tools for efficiency and shift workforce focus toward collecting contextual information directly from suppliers.

### **New Entrants: Emphasis vs. Outcome**

While attrition data points to sustained erosion, the DoW's response has focused less on reducing supplier outflows than on attracting new participants. In 2014, it announced a strategic imperative to draw innovative commercial companies into the defense market (Hummel and Wurster). This shift reflected a growing recognition that commercial innovation with defense applications was outpacing the traditional DIB. In 2020, the COVID pandemic exposed the DoW to latent supply chain vulnerabilities and the consequences of decades of consolidation and off-shoring.

Taken together, these conditions prompted the DoW to direct billions over the last decade into new and legacy initiatives touted as opportunities to revitalize the DIB. Each of these programs shares a common premise: simplifying aspects of the defense procurement process will attract innovative commercial companies into the defense market. If effective, then, these strategies should result in more first-time vendors entering the DIB every year. However, as shown in Figure 9, the count of new vendors annually has remained stable at best. Additionally, as shown in Figure 10, new entrants' share of contract funding has remained economically marginal.

Count and share of annual vendor pool new to DoW

| FY   | Total Vendors | New to DoW | % New |
|------|---------------|------------|-------|
| 2015 | 67,797        | 7,886      | 11.6% |
| 2016 | 65,340        | 7,426      | 11.4% |
| 2017 | 63,537        | 7,003      | 11.0% |
| 2018 | 60,419        | 6,367      | 10.5% |
| 2019 | 57,374        | 5,765      | 10.0% |
| 2020 | 53,965        | 5,698      | 10.6% |
| 2021 | 52,281        | 5,146      | 9.8%  |
| 2022 | 48,700        | 4,094      | 8.4%  |
| 2023 | 48,554        | 4,721      | 9.7%  |
| 2024 | 48,035        | 5,080      | 10.6% |

**Figure 9. New Vendors Entering the Defense Market**





Share of defense spending going to new-to-DoW entrants

| FY   | Total DoW | New Vendor \$ | % to New |
|------|-----------|---------------|----------|
| 2015 | \$274.7B  | \$3.1B        | 1.1%     |
| 2016 | \$295.7B  | \$3.1B        | 1.1%     |
| 2017 | \$323.7B  | \$4.0B        | 1.2%     |
| 2018 | \$357.9B  | \$5.3B        | 1.5%     |
| 2019 | \$393.9B  | \$9.2B        | 2.3%     |
| 2020 | \$436.2B  | \$6.3B        | 1.4%     |
| 2021 | \$400.0B  | \$11.9B       | 3.0%     |
| 2022 | \$422.4B  | \$9.8B        | 2.3%     |
| 2023 | \$472.3B  | \$4.6B        | 1.0%     |
| 2024 | \$461.7B  | \$6.0B        | 1.3%     |

**Figure 10. Obligations to New Defense Vendors**

### New Vendor Inflation via Affiliation

Our methodology treats each first-time UEI as a new vendor, without accounting for instances in which the “new” entity actually represents a joint venture or teaming arrangement between existing defense contractors. To assess the potential impact of these affiliations on our count of new entrants, we built a process to identify UEIs classified as “new” that likely represent a reconstitution of one or more existing DoW contractors.

First, we built a logic chain to detect potential joint ventures (JVs), parent-subsidary relationships, and instances in which multiple vendors share the same domain name. We analyzed vendor names to identify mentions of “Joint Venture,” “JV,” “Consortium,” or “Teaming,” excluding instances where these terms may have a different association. For instance, the process identifies “AECOM-Baker JV” and “3T Federal-BBIX JV LLC” as joint ventures, but rejects “JVS Engineering,” as “JVS” is not assumed to be an abbreviation of joint venture. We extracted the constituent party names from joint venture vendor names to identify the underlying entities. Next we examined the FPDS field for “parent UEI” to detect potential subsidiaries of larger entities. For new entrants with an associated parent UEI, we flagged potential subsidiary relationships where the parent entity differed from the new entrant. Lastly we looked at [SAM.gov](https://sam.gov) registration data via company CAGE codes to detect if two different vendors shared a company website. If a “new” vendor had the same website as an existing defense contractor—excluding common hotel chains and other false positives—we implied a common connection. As shown in Figure 11, approximately 14% of companies we identified as new entrants had an affiliation signal, suggesting they are likely extensions of existing DoW contractors.



| Type                            | Count  | % of New-to-DoW |
|---------------------------------|--------|-----------------|
| Joint Venture (name match)      | 1,579  | 2.7%            |
| Subsidiary (parent UEI differs) | 5,849  | 9.9%            |
| Shared Website Domain Only      | 675    | 1.1%            |
| Any Affiliation Signal          | 8,103  | 13.7%           |
| Presumed Independent            | 51,083 | 86.3%           |

**Figure 11. Affiliation Signals Among New Defense Entrants**

The presence of affiliations in the new vendor data reflects a larger DIB trend: entrenched entities find ways to expand their market share over time, and spinoffs/partnerships are one such mechanism. For example, partnering with small, woman-owned, or disadvantaged businesses may enable them to qualify for contracting opportunities otherwise unavailable to them. They are not abusing the system; they are using it to their advantage. But to the extent new suppliers are simply amalgamations of existing players, the objectives of industrial base growth and diversification have not been met—the DoW continues to depend on the same limited pool of suppliers, and innovative commercial companies remain out of reach. When examining how policies affect the composition of the DIB, the DoW must flag affiliated entities for further analysis.

### Who Entered?

Given the DoW's stated priority of attracting innovative companies into the DIB, we next sought to examine whether the companies that entered the defense market reflected that ambition. Although qualifying a company as "innovative" is highly subjective, we leveraged available data to develop a meaningful proxy. We assigned an "innovation score," from 0-10, to each new vendor using objective signals from their entry-year DoW contract actions:

- DoW SBIR/STTR contract action(s): +3 points
- Other Transaction Agreement (OTA) prototype: +2 points
- OTA production: +1 point
- R&D PSC: +2
- Commercial product or service flag: +1
- Small business designation: +1
- IT/Software PSC: +1

A vendor scoring 7+ had multiple strong innovation signals; 3-6 reflected a moderate score; and 0-2 signaled a traditional/non-innovative supplier. Using this scoring methodology, we found that 83% of new entrants during our analysis period were "traditional," rather than innovative suppliers. Figure 12 provides a breakdown of new entrants by innovation score annually. These results are a damning indictment of more than a decade of policy explicitly aimed at creating pathways for innovative commercial companies to break into the defense market.






**Figure 12. Innovativeness Among New Vendors, Annually**

## Innovation Programs

New entrants received higher points for leveraging SBIR/STTR and OTAs because both share a stated purpose of stimulating defense-focused R&D and enabling commercial innovators to engage with the DoW without the burdens of the traditional contracting process. In other words, they serve as a theoretical marker for innovativeness. Given that so few new entrants scored highly suggested a low incidence of SBIR/STTR and OTAs among them. For a closer look, we isolated each new entrants' initial contracting action and, using FPDS data, categorized it as a SBIR/STTR, Other Transaction (OT), or traditional contract. As shown in Figure 13, the vast majority of new vendors enter the DIB via traditional procurement contracts.

OTs, which would include DIU—another innovation program marketed as a way for commercial innovators to break into the defense market—account for a negligible share of new entrants' initial contracts. For the majority of our analysis period, fewer than 10% of new entrants leveraged DoW SBIR/STTR to break into the defense market in a given year. Gains in new entrants via SBIR/STTR between FY2019-FY2021 likely attribute to the launch of “Open Topics” in the Air Force SBIR/STTR program. The Open Topic model allows companies to submit a simplified proposal that explains how their solution could solve a defense-related problem, as opposed to the traditional SBIR/STTR process which requires companies to respond to prescriptive topics. The subsequent reversal in new vendors suggests that these mechanisms either lost efficacy or became another tool for incumbents to expand their market share. In short, most new entrants are not innovative companies, and initiatives marketed as on-ramps for innovative commercial firms play a minor role in expanding the industrial base.

First procurement action, FY2015–FY2024. Categories are mutually exclusive.

| FY   | New Vendors | Contract      | SBIR/STTR   | OTA       |
|------|-------------|---------------|-------------|-----------|
| 2015 | 7,886       | 7,748 (98.2%) | 133 (1.7%)  | 5 (0.1%)  |
| 2016 | 7,426       | 7,286 (98.1%) | 130 (1.8%)  | 10 (0.1%) |
| 2017 | 7,003       | 6,803 (97.1%) | 157 (2.2%)  | 43 (0.6%) |
| 2018 | 6,367       | 6,194 (97.3%) | 143 (2.2%)  | 30 (0.5%) |
| 2019 | 5,765       | 5,168 (89.6%) | 512 (8.9%)  | 85 (1.5%) |
| 2020 | 5,698       | 5,006 (87.9%) | 593 (10.4%) | 99 (1.7%) |
| 2021 | 5,146       | 4,559 (88.6%) | 511 (9.9%)  | 76 (1.5%) |
| 2022 | 4,094       | 3,721 (90.9%) | 303 (7.4%)  | 70 (1.7%) |
| 2023 | 4,721       | 4,268 (90.4%) | 367 (7.8%)  | 86 (1.8%) |
| 2024 | 5,080       | 4,587 (90.3%) | 410 (8.1%)  | 83 (1.6%) |

**Figure 13. Entry Path of New Defense Vendors**

### New Entrants & Retention

Assessing the effectiveness of policies designed to expand the industrial base requires not only evaluating how well they attract new vendors, but also how new entrants perform in the DIB over time. To that end, we next looked at the tenure of new entrants in the defense market. We joined new entrants’ contract actions for the entirety of our analysis period and isolated the most recent DoW funding action for each firm. We then bucketed new entrants into six categories:

- 1) Companies that lasted in the defense market for less than one year (they had their first and last DoW contract action in the same 12-month period)
- 2) New entrants that lasted one year in the defense market, then left
- 3) New entrants that lasted two years in the defense market, then left
- 4) New entrants that worked with DoW for 3-5 years, then left
- 5) New entrants that worked with DoW for 6+ years, then left
- 6) New entrants that remained active in the defense market as of FY2024

As shown in Figure 14, nearly 1/3 of new entrants survived less than 12 months in the defense market, and more than 43% departed after one year. Based on these findings, a significant share of new vendors do not become meaningful contributors to the DIB.

Analyzing PSCs for single-year vendors, we found that approximately 15% associated with hotels for overnight travel—instances where the DoW stayed at a new hotel once during a 12-month period, and did not return. The remaining single-year vendors cross-cut industries, including R&D, professional services, and hardware. Further research is needed to understand why so many new vendors exit the defense market within their first year. However, DoW leaders and policymakers often lament the “Valley of Death”—a byproduct of bureaucracy in the defense acquisition system that prevents companies from successfully converting initial R&D/prototype funding into DoW programs of record (Landreth). Our analysis of new entrants’ tenure provides



an empirical framework for calculating the Valley of Death in real terms. It also reinforces the need for the DoW to improve efforts to both attract innovative commercial technologies into the DIB, as well as to retain—and maximally leverage—the most promising among them.

Figure 14: Tenure of New Entrants in Defense Market

| New Entrant Status                                    | Count, New Entrants | % Total New Entrants |
|-------------------------------------------------------|---------------------|----------------------|
| Entered & Exited in the Same Year (Active <12 months) | 18,612              | 31%                  |
| Exited After 1 Year                                   | 6,835               | 12%                  |
| Survived 2 Years Then Left                            | 3,768               | 6%                   |
| Survived 3–5 Years Then Left                          | 6,506               | 11%                  |
| Survived 6+ Years Then Left                           | 1,960               | 3%                   |
| Active as of 2024                                     | 21,508              | 36%                  |

*Figure 14. Tenure of New Entrants in Defense Market*

## Conclusions & Recommendations

The DIB is in crisis. With each passing year, defense spending skyrockets while the contractor pool shrinks. For every new company that enters the defense market, two companies leave. Given the range of products and services affected by attrition, these trends are less likely the result of strategic consolidation and more likely a consequence of rising compliance costs, slow acquisition cycles, and budget uncertainty. As these regulatory burdens force smaller firms out of the DIB, they enable the largest defense contractors—that can absorb them easily—to expand their market share. The concentration that follows leaves the DoW dangerously exposed, with few alternatives if and when a critical supplier faces production issues or raises national security concerns.

Despite billions in funding for industrial base expansion, the DoW has achieved no meaningful net growth in new entrants over the past decade. Innovation initiatives do not attract new vendors into the defense market, and the majority of first-time DIB suppliers are not innovative commercial companies. Efforts to support transition have likewise failed: more than 40% of new entrants leave the defense market within one year.

The health of the industrial base is a matter of national security, and the DoW must treat the collapsing DIB as a real-time threat. The Office of the Secretary of War (OSW) must maintain a real-time common operating picture of the DIB, continuously monitoring new entrants, attrition, and market consolidation. Bureaucratic fluency is not a warfighting advantage, thus the DoW must require agencies to meet sector-specific goals for contracting with new and existing suppliers. DoW personnel should engage with suppliers directly, so that impediments to achieving these goals can be identified and mitigated immediately.

The current trends are reversible—but only with new incentives, a deliberate reallocation of resources, and committed leadership willing to resist the lobbying campaigns of large defense contractors.





## Reference List

AAR CORP. "AAR Awarded First Captains of Industry Contract by the U.S. Department of Defense's DLA Focused on Distribution of Sustainment Services." *AAR CORP.*, 19 July 2022, [www.aarcorp.com/en/newsroom/press-releases/2022/aar-awarded-first-captains-of-industry-contract-by-the-u.s.-department-of-defenses-dla-focused-on-distribution-of-sustainment-services](http://www.aarcorp.com/en/newsroom/press-releases/2022/aar-awarded-first-captains-of-industry-contract-by-the-u.s.-department-of-defenses-dla-focused-on-distribution-of-sustainment-services).

Hummel, Robert, and Kathryn Schiller Wurster. "Department of Defense's Innovation Experiment." *STEPS: Science, Technology, and Engineering Policy Studies*, Potomac Institute for Policy Studies, 30 June 2016, [www.potomac institute.org/steps/index.php/issues/department-of-defense-s-innovation-experiment](http://www.potomac institute.org/steps/index.php/issues/department-of-defense-s-innovation-experiment).

Landreth, James M. "Through DoD's Valley of Death—A Data-Intensive Startup's Journey." *Defense Acquisition*, Defense Acquisition University, Jan.–Feb. 2022, [www.dau.edu/library/defense-atl/blog/Valley-of-Death](http://www.dau.edu/library/defense-atl/blog/Valley-of-Death).

Serbu, Jared. "Defense Industrial Base Loses Thousands More Firms, Mainly Because DoD Is a Tough Customer." *Federal News Network*, 10 Feb. 2023, [federalnewsnetwork.com/defense-industry/2023/02/defense-industrial-base-loses-thousands-more-firms-mainly-because-dod-is-a-tough-customer/](http://federalnewsnetwork.com/defense-industry/2023/02/defense-industrial-base-loses-thousands-more-firms-mainly-because-dod-is-a-tough-customer/).

U.S. Census Bureau. "North American Industry Classification System (NAICS)." *United States Census Bureau*, [www.census.gov/naics/](http://www.census.gov/naics/). Accessed 25 Mar. 2026.



# Buying Flexibility: Converting Cost-Plus to Firm-Fixed-Price to Strengthen Mission Resilience

**M. Jay Barlow**—is a senior acquisition professional with extensive experience in defense contracting, contract-type strategy, and cost analysis across national security space programs. His work focuses on the intersection of contract structure, risk allocation, and mission outcomes, with particular emphasis on optimizing the government's approach to cost-reimbursement and fixed-price instruments. A retired Navy Supply Corps officer with more than 20 years of uniformed service—including contingency operations in Afghanistan—he brings more than 30 years of experience in defense acquisition, contracting strategy, and capability delivery. He holds advanced credentials in defense acquisition and continues to research practical frameworks for improving acquisition agility and warfighter readiness.  
[mbarlow7@capellauniversity.edu]

## Abstract

The Department of War has long relied on cost-reimbursement (CP) contracts to manage technical uncertainty, but continued use of cost-type instruments on mature, stable efforts imposes unnecessary administrative burden, diffuses contractor accountability, and slows delivery of warfighting capability. This paper proposes a readiness-gate framework for determining when and how to convert CP contracts to firm-fixed-price (FFP) arrangements. Drawing on Federal Acquisition Regulation authorities, historical case studies from space and defense programs, Government Accountability Office assessments, and recent senior-leader direction to accept greater acquisition risk in order to reduce operational risk, the study defines eight readiness gates—including requirement maturity, configuration stability, cost-data transparency, and supply-chain robustness—and operationalizes them as go/no-go criteria. It develops five archetypal conversion pathways, ranging from progressive risk transfer through fixed-price incentive and economic-price-adjustment instruments to direct CP-to-FFP conversion for mature recurring work. A practitioner toolkit provides checklists, sample contract line-item structures, financing-method comparisons, and risk-benefit tradeoff analysis. The paper argues that disciplined, selective conversion improves cost discipline, reduces audit and compliance burden, and frees resources for mission-focused oversight—while cautioning that premature FFP conversion can increase price, default risk, and change-order costs.

**Keywords:** contract-type conversion, firm-fixed-price, cost-reimbursement, acquisition risk, mission resilience, defense industrial base

## Introduction

Contract-type selection is one of the most consequential decisions in defense acquisition. It determines how cost risk is allocated between the government and the contractor, shapes the incentive environment for performance, and defines the administrative burden that both parties will bear throughout execution. The Federal Acquisition Regulation (FAR) establishes a spectrum of contract types ranging from cost-plus-fixed-fee (CPFF), which places virtually all cost risk on the government, to firm-fixed-price (FFP), which transfers maximum cost risk to the contractor (FAR 16.103). Between these poles lie instruments such as cost-plus-incentive-fee (CPIF), fixed-price incentive firm (FPIF), and fixed-price with economic price adjustment (FPEPA) that share risk in varying proportions.

Despite decades of reform—from the Weapon Systems Acquisition Reform Act of 2009 (WSARA; Pub. L. 111-23), which mandated competitive prototyping and independent cost estimates to improve acquisition outcomes, to the Better Buying Power initiatives that encouraged greater use of fixed-price instruments—the Department of War (DoW) continues to rely heavily on cost-reimbursement contracts well into system maturity, even when technical baselines have stabilized and production processes are well understood (Hastings et al., 2016). A 2021 Government Accountability Office (GAO) review found that while the DoW's use of fixed-



price incentive (FPI) contracts on major defense acquisition programs (MDAPs) had increased from 17% to 49% of total obligations between fiscal years 2010 and 2019, the department had not systematically assessed whether this shift actually contributed to better cost and schedule outcomes (GAO, 2021). This pattern persists for understandable reasons: cost-type contracts reduce the contracting officer's immediate risk of protest or default, provide the government with extensive cost visibility, and accommodate the requirements changes that characterize complex weapon-system development. However, these benefits come at significant cost. Cost-reimbursement instruments require approved accounting systems, generate intensive incurred-cost submissions, demand audit resources from the Defense Contract Audit Agency (DCAA), and—perhaps most critically—dilute contractor incentives to control costs and optimize production processes. MacDonald (2024) observes that cost-plus contracting limits financial upside for contractors while the defense acquisition process creates incentive misalignments that prioritize compliance over value creation.

Recent senior leadership has called for a fundamental shift in acquisition culture. In November 2025, the Secretary of War issued directives to transform the defense acquisition system into a “warfighting acquisition system,” explicitly calling for greater institutional tolerance for acquisition risk in order to reduce operational risk and accelerate fielding of urgently needed capabilities (Secretary of War, 2025). This direction frames contract-type conversion not as an academic exercise but as an operational imperative: where technical and market conditions permit, shifting from cost-type to fixed-price contracts can free obligation authority, reduce administrative drag, and incentivize the industrial base to invest in efficiency rather than relying on cost reimbursement as a hedge.

This paper responds to that imperative by proposing a practical, decision-oriented framework for converting cost-reimbursement efforts to firm-fixed-price arrangements. It argues that disciplined CP-to-FFP conversion should be treated not as a blanket ideology but as a selective, readiness-based acquisition decision. When requirements, configuration, cost data, and supply-chain conditions are sufficiently mature, FFP can improve cost discipline, mission resilience, and administrative efficiency while rewarding industry for accepting performance risk. When those conditions are not met, the government should retain cost-type instruments or use bridging mechanisms such as FPIF or FPEPA rather than forcing premature fixed pricing.

## **Why CP-to-FFP Conversion Matters Now**

Three converging pressures make deliberate contract-type conversion a timely policy concern. First, the fiscal environment demands greater budget predictability. Cost-reimbursement contracts expose the government to open-ended financial liability; actual costs may exceed estimates, requiring reprogramming actions that disrupt portfolio planning and delay other programs. The GAO's 2024 annual assessment of 108 selected weapon programs found that DoW plans to invest more than \$2 trillion to develop and acquire its costliest systems, yet nearly half of the 31 major programs assessed in depth reported cost increases in the past year, and the average time to deliver initial capability had grown from roughly eight years to eleven years—an average increase of three years from original planned dates (GAO, 2024). While contract type alone does not explain this growth, open-ended cost-reimbursement arrangements contribute to a fiscal environment in which overruns are accommodated rather than constrained. FFP contracts, by contrast, establish a ceiling on the government's financial obligation for defined scope, enabling more reliable obligation-authority management across mission threads (Barlow, 2026).

Second, the administrative burden of cost-type contracting consumes scarce government and industry talent. Under cost-reimbursement arrangements, contractors must maintain approved cost-accounting systems, submit detailed incurred-cost proposals, and



undergo extensive DCAA audits. The government, in turn, must staff program offices with cost analysts, auditors, and financial managers to review these submissions. The Congressional Research Service has noted that the DoW's contract-pricing decisions directly affect both the efficiency of defense spending and the health of the industrial base, with contract type serving as the primary lever for allocating cost risk and shaping contractor behavior (Peters, 2024). When a program transitions to FFP, incurred-cost submissions are eliminated, earned-value management (EVM) clauses may be removed unless the risk profile warrants retention, and oversight shifts from cost tracking to delivery, quality, and acceptance metrics (Barlow, 2026). This reallocation of talent from compliance to mission-focused oversight represents a significant efficiency gain.

Third, and most important, is the warfighting imperative. Senior leaders have articulated a clear expectation that the acquisition enterprise move from bureaucratic risk aversion toward faster delivery of operationally relevant capability. The Office of the Secretary of Defense for Defense Pricing and Contracting (then DPAP) found in 2016 CPIF and FPIF contracts were correlated with better cost and schedule outcomes than either pure CPFF or FFP arrangements, and explicitly cautioned against protracted use of cost-reimbursement instruments when technical maturity would support incentive-based structures (Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, 2016). The GAO's separate review of cost-type contract usage found that contract type alone does not determine cost or schedule outcomes—knowledge-based acquisition practices and well-structured incentives matter more than the contract-type label (GAO, 2020). Together, these findings support the view that where technical maturity justifies it, shifting risk to the contractor through fixed-price mechanisms can sharpen incentives for efficiency without degrading mission success, provided the conversion is based on sound readiness criteria rather than ideological preference.

Importantly, the Institute for Defense Analyses (IDA) found no statistically significant relationship between contract type and cost growth in a comprehensive analysis of defense programs (Arnold et al., 2008). MacDonald (2024) reinforces this conclusion through a financial lens, demonstrating that compliance-focused acquisition processes erode the value of defense spending regardless of contract type, and that the real question is how to structure incentives that maximize capability returns per dollar invested. These findings reinforce the central argument of this paper: the question is not whether to use FFP, but when and how to transition to it responsibly.

## **Readiness Gates for Conversion**

FAR 16.104 enumerates the factors a contracting officer must consider when selecting a contract type, including price competition, price and cost analysis, type and complexity of requirement, urgency, period of performance, contractor capability, adequacy of accounting systems, and acquisition history. FAR 16.202-1 further specifies that FFP is appropriate when the contracting officer can establish fair and reasonable prices at the outset based on reasonably definite functional or detailed specifications. The Section 809 Panel's Advisory Panel on Streamlining and Codifying Acquisition Regulations recommended simplifying the contract-type selection process and reducing regulatory barriers to efficient contracting, reinforcing the need for clear, practical decision criteria (Section 809 Panel, 2019). Building on these regulatory foundations and the empirical findings of the Hastings et al. (2016) acquisition-efficiency study, this paper defines eight readiness gates that acquisition teams should evaluate before initiating a CP-to-FFP conversion.

### **Gate 1: Requirement Maturity**

Requirements must be clear, stable, and realistic. As Hastings et al. (2016) emphasize, requirements stability is essential for fixed-price contracting because costs cannot be controlled



when the contract is opened for changes. The Weapon Systems Acquisition Reform Act of 2009 codified the principle that major defense programs should demonstrate a high level of knowledge—including mature requirements and validated technology—before proceeding past critical decision points (Pub. L. 111-23). A program should demonstrate that its key performance parameters and key system attributes have been baselined, validated by the requirements community, and stable for a sufficient period—typically two or more production cycles or 12 months of sustainment—before conversion is considered.

### **Gate 2: Configuration Stability**

Configuration control must be established with frozen external interfaces. Engineering change proposals (ECPs) should be at a low, predictable rate. Historical experience across defense programs demonstrates that frequent ECPs under fixed-price contracts generate significant friction, as contractors interpret scope narrowly and government-induced changes become inevitable. Configuration stability reduces this risk by ensuring that the technical baseline on which the fixed price rests is unlikely to shift materially during execution.

### **Gate 3: Interface Stability**

For systems integrated into larger architectures, external interfaces must be formally controlled and unlikely to change. Unstable interfaces—particularly in multi-segment space systems or software-intensive platforms—can drive unpredictable integration costs that are difficult to price under FFP. Where interface risks remain, a hybrid CLIN structure with a cost-type CLIN for integration engineering may be appropriate.

### **Gate 4: Technical and Integration Risk**

Technical risk must be minimal. FAR 16.202 implicitly requires this by limiting FFP to efforts with a firm pricing basis. The Hastings et al. (2016) study identifies mature technology as a prerequisite, noting that technology readiness level (TRL) alone is insufficient—manufacturing readiness, software complexity, and integration risk must also be assessed. Congress reinforced this principle in Section 808 of the FY2023 National Defense Authorization Act, which limits the number of low-rate initial production (LRIP) lots to no more than one when the milestone decision authority authorizes a fixed-price contract at Milestone B and the scope includes both development and LRIP. The service acquisition executive may waive this limitation, but must notify the congressional defense committees within 30 days (Pub. L. 117-263, § 808; DFARS Case 2023-D009). Programs at TRL 5–6 with unresolved integration challenges, should generally receive a no-go determination for FFP.

### **Gate 5: Production and Process Repeatability**

The effort must demonstrate repeatable processes and predictable cycle times. For production programs, this means stable bills of material, demonstrated manufacturing yields, and consistent labor-hour expenditures across recent lots. For sustainment efforts, this means predictable workload volumes and stable key performance indicators over an extended period. The Barlow (2026) framework suggests a threshold of cost and schedule variance within plus or minus 3% over the most recent three months for production programs, or plus or minus 1.5% over 24 months for sustainment efforts.

### **Gate 6: Cost-Data Transparency and Pricing Basis**

A firm basis for pricing must exist. FAR 15.403-4 requires certified cost or pricing data for negotiated contract actions above applicable thresholds unless an exception applies. The Congressional Research Service has highlighted that the adequacy of the government's pricing data—including access to actual costs, competitive market intelligence, and historical performance data—is a critical determinant of whether fixed-price arrangements will yield favorable outcomes (Peters, 2024). For conversion purposes, the government should have





access to recent actual costs, credible bases of estimate (BOEs), and sufficient cost history to support price analysis or cost realism assessment. Where cost-type contracts have generated actual-cost data over multiple performance periods, this data provides the pricing foundation that FFP demands. As the Hastings et al. (2016) study notes, historical experience indicates that a firm pricing basis generally does not exist until at least one product unit has been delivered successfully.

### **Gate 7: Supply-Chain Robustness and Volatility**

Supply-chain stability must be evaluated, including supplier availability, lead times, commodity price volatility, and labor market conditions. Where inflation or material-cost indices are volatile, FPEPA (authorized under FAR 16.203) may be more appropriate than pure FFP. The contracting officer should assess whether key suppliers are on firm quotes, whether long-lead parts have been secured, and whether the material mix is insensitive to commodity shocks. If multi-year price stability is a concern, FPEPA with adjustment clauses tied to published indices can protect both parties while still transferring production risk to the contractor.

### **Gate 8: Funding Stability**

Fixed-price contracts require funding stability through the life of the contract to meet contractually established payment obligations (Hastings et al., 2016). Budget instability—including continuing resolutions, sequestration risk, or anticipated program restructuring—can undermine FFP execution by forcing the government to delay orders, reduce quantities, or renegotiate terms, all of which erode the pricing basis. The acquisition team should confirm that the program's funding profile aligns with the proposed contract's payment schedule and that multi-year procurement or economic-order-quantity authorities are available where appropriate.

### **Go/No-Go Decision Criteria**

Table 1 translates the eight readiness gates into explicit decision criteria. A “Go” determination requires that all critical gates are satisfied. A “Conditional Go” indicates that the program may proceed with conversion using a bridging instrument or hybrid CLIN structure to manage residual risk. A “No-Go” determination means the program should retain cost-type contracting and establish measurable exit criteria for future reassessment.



**Table 1. Readiness-Gate Go/No-Go Decision Criteria**

| Readiness Gate             | Go                                                | Conditional Go                                                   | No-Go                                                   |
|----------------------------|---------------------------------------------------|------------------------------------------------------------------|---------------------------------------------------------|
| Requirement Maturity       | KPPs/KSAs baselined and stable >12 months         | Minor requirements TBDs remain; can be isolated in separate CLIN | Active requirements churn; user community uncommitted   |
| Configuration Stability    | Config controlled; ECP rate low and predictable   | Moderate ECP rate; manageable through change-order CLIN          | Frequent design changes; configuration not yet frozen   |
| Interface Stability        | Interfaces formally controlled and frozen         | Most interfaces stable; residual risk in one segment             | Multiple undefined or evolving interfaces               |
| Technical/Integration Risk | TRL 8–9; integration demonstrated                 | TRL 7; integration underway with known path                      | TRL 5–6; unresolved known unknowns                      |
| Process Repeatability      | Stable BOMs; variance $\leq \pm 3\%$              | First units complete; learning curve predictable                 | No production history; first-article pending            |
| Cost-Data Transparency     | Actual cost history from $\geq 2$ lots or periods | BOEs supported by analogous program data                         | No credible pricing basis; estimates only               |
| Supply-Chain Robustness    | Firm supplier quotes; stable lead times           | Moderate volatility; manageable via FPEPA                        | Critical sole-source suppliers; commodity shocks likely |
| Funding Stability          | Funded through contract period; MYP if applicable | CR environment but stable topline expected                       | Major budget uncertainty; restructure likely            |

## Conversion Pathways and Selection Logic

Not every CP-to-FFP conversion follows the same path. The appropriate pathway depends on the program’s maturity, risk profile, and the specific gates that remain open or partially satisfied. This paper identifies five archetypal conversion pathways, each mapped to the conditions under which it is most suitable.

### Pathway 1: Progressive Risk Transfer (CP → FPIF → FFP)

This pathway is appropriate for programs transitioning from development to production where residual cost uncertainty exists but is declining. The program begins with cost-type development, moves to FPIF for initial production lots (sharing cost risk through target costs, ceilings, and share ratios per FAR 16.204), and converts to FFP once production processes stabilize and a firm pricing basis is established. The 2016 OSD guidance on incentive contracts found that FPIF arrangements were correlated with better cost outcomes than either pure cost-plus or pure fixed-price instruments during transitional production phases, suggesting that progressive risk transfer through incentive-fee structures can serve as an effective bridge to full FFP (Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, 2016). The GAO confirmed this trend, reporting that the DoW’s use of FPI contracts on MDAPs grew from 17% of obligations in fiscal year 2010 to 49% by fiscal year 2019, though the department had not formally assessed whether this increased usage contributed to improved program outcomes (GAO, 2021). This progressive approach aligns with MacDonald’s (2024) finding that the defense acquisition system’s high transaction costs and compliance burden create an environment where incremental risk transfer—rather than abrupt contract-type changes—better



serves both government and industry interests. FPIF contracts require an approved accounting system and provide the government with earned-value and cost data that can anchor future FFP negotiations.

## **Pathway 2: CP Development with FPEPA Production**

Where production economics are well understood but multi-year commodity or labor volatility creates pricing risk, the program can transition from CP development directly to FFP with economic price adjustment clauses tied to published indices (FAR 16.203). This pathway is particularly suitable for sustainment and operations-and-maintenance efforts where workload is predictable but inflationary pressures are real. The ORION-O&M scenario in the Barlow (2026) framework illustrates this application: a five-year CP sustainment contract with stable KPIs and predictable demand is a strong candidate for FFP conversion, with FPEPA clauses protecting against wage growth that, while modest at approximately 3% annually, compounds over a multi-year performance period.

## **Pathway 3: Hybrid CLIN Model**

Many programs contain a mix of mature, recurring work and uncertain engineering or integration tasks. Rather than forcing the entire effort into a single contract type, the hybrid CLIN model retains cost-type CLINs for uncertain work—such as engineering change orders, integration of government-furnished equipment, or unresolved interface engineering—while converting stable, recurring CLINs to FFP. This approach supports risk segmentation within a single contract and avoids the all-or-nothing dynamic that can lead to either premature FFP (with excessive risk) or unnecessary CP retention (with insufficient incentives). Programs that segregate uncertain engineering work from stable production work have historically achieved better outcomes by allowing each CLIN to carry risk-sharing mechanisms appropriate to its maturity level.

## **Pathway 4: Direct CP-to-FFP Conversion**

For mature, recurring sustainment or production efforts that satisfy all eight readiness gates, direct conversion from cost-type to FFP is appropriate. This pathway applies when the program has extensive cost history, demonstrated process repeatability, stable requirements, and a robust supply chain. The conversion may be executed via bilateral supplemental agreement (SF-30) per FAR Part 43 if the change is within scope, or through a new competitive acquisition if the scope determination reveals a material change in deliverables, dollar value, duration, or risk—the “4 Ds + C” test for scope integrity (Barlow, 2026).

## **Pathway 5: No-Go (Retain CP)**

Where readiness gates are not satisfied, the correct answer is to retain cost-type contracting and establish measurable exit criteria for later reassessment. Exit criteria should be specific, time-bound, and tied to observable program events—such as completion of engineering-model testing, achievement of interface control baseline, cost and schedule variance within defined thresholds for a specified number of builds, or demonstration of manufacturing repeatability. This pathway is not a failure; it is a disciplined recognition that premature FFP conversion can increase price (because the contractor will embed larger risk margins), elevate default risk, generate costly change orders, and reduce the government's visibility into actual program costs (Hastings et al., 2016). The GAO's review of Navy shipbuilding found that FPI contracts on lead ships with immature designs frequently resulted in cost overruns that exceeded ceiling prices, forcing the government to absorb losses through follow-on contract adjustments and eroding the incentive structure that fixed pricing was intended to create (GAO, 2017). MacDonald (2024) provides a broader fiscal context for why premature fixed pricing is dangerous: defense contractors already face unique bureaucratically imposed financial risks, steep discounting of future cash flows, and limited upside under cost-



plus arrangements. When fixed-price contracts are imposed prematurely, these pressures compound, driving contractors to either embed excessive risk margins or decline to participate—a trend that further weakens the defense industrial base.

## Practitioner Toolkit for Structuring Conversions

### Readiness-Gate Checklist

The checklist in Table 2 provides acquisition teams with a structured assessment instrument. Each gate should be evaluated with supporting evidence documented in the acquisition strategy or contract-type determination memorandum. The checklist is designed to be iterative: programs that receive a no-go or conditional-go determination should schedule reassessment at defined program milestones.

**Table 2. Readiness-Gate Assessment Checklist**

| Gate                          | Key Evidence Required                                                               | Assessment               |
|-------------------------------|-------------------------------------------------------------------------------------|--------------------------|
| 1. Requirement Maturity       | Approved CDD/CPD; KPP stability period documented; user sign-off                    | Go / Conditional / No-Go |
| 2. Configuration Stability    | Configuration audit results; ECP rate trend; interface control document status      | Go / Conditional / No-Go |
| 3. Interface Stability        | ICD approval status; external dependency register; segment integration test results | Go / Conditional / No-Go |
| 4. Technical/Integration Risk | TRL/MRL assessment; integration test results; risk register burn-down               | Go / Conditional / No-Go |
| 5. Process Repeatability      | Unit cost actuals for ≥2 lots; cycle time data; yield/rework rates                  | Go / Conditional / No-Go |
| 6. Cost-Data Transparency     | Actual cost history; BOE/BOM validation; DCAA audit findings                        | Go / Conditional / No-Go |
| 7. Supply-Chain Robustness    | Supplier health assessment; material price indices; lead-time trends                | Go / Conditional / No-Go |
| 8. Funding Stability          | POM alignment; CR/sequestration risk assessment; MYP/ESP eligibility                | Go / Conditional / No-Go |

### Scope-Integrity Decision Checklist

Before executing a contract-type conversion via modification, the contracting officer must determine whether the conversion constitutes an in-scope change. Improper conversions risk violating the Competition in Contracting Act (CICA). The scope-of-the-competition test asks whether original offerors would have reasonably anticipated the change. The Barlow (2026) framework operationalizes this through the 4 Ds + C assessment: Deliverables—are the items fundamentally the same? Dollars—does the total value change materially? Duration—does the period of performance change significantly? Degree of risk—does the risk profile shift enough to alter competitive dynamics? Competition—would the change have attracted different offerors? If changes are minor and expected, the conversion may proceed via bilateral supplemental agreement using SF-30 per FAR Part 43. If changes are material, the government must recompet or obtain a justification and approval (J&A) under FAR Part 6.

### Sample CLIN/SLIN Structuring Patterns

Effective CLIN structuring is critical to successful conversion. Table 3 presents three representative patterns that accommodate different risk profiles.



**Table 3. Sample CLIN Structuring Patterns for Conversion**

| Pattern         | CLIN Structure                                                                                    | Suitable Conditions                                                       |
|-----------------|---------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
| Full Conversion | All CLINs convert from CPFF/CPIF to FFP                                                           | All gates satisfied; mature recurring work; stable scope                  |
| Hybrid          | FFP for production/sustainment CLINs; CP for engineering-change or integration CLINs              | Production mature but residual development tail; uncertain interface work |
| Phased          | Initial lots under FPIF; later lots convert to FFP as confidence grows; CP SLIN for emergent work | Early production with declining but nonzero risk; learning-curve pricing  |

### **Financing Options: Progress Payments Versus Performance-Based Payments**

The choice of financing method is a critical component of FFP conversion. FAR 32.1001 establishes performance-based payments (PBPs) as the preferred government financing method. Under PBPs, payments are tied to objective, quantifiable measures of performance—milestones, test completions, or deliverable acceptances—rather than incurred costs (FAR 32.1004). This alignment between payment and performance reinforces the incentive structure of FFP contracts.

Progress payments (FAR 52.232-16), by contrast, reimburse a percentage of incurred costs—80% for large businesses and 85% for small businesses. While simpler to administer, progress payments retain a cost-reimbursement flavor that partially undermines the cost-discipline advantages of FFP. Importantly, progress payments and PBPs may not be used simultaneously on the same effort or CLIN. The acquisition team must select one method and ensure that the payment terms align with CLIN/Sub-CLIN structure and DFAS disbursement instructions to avoid payment throttling.

For conversion scenarios, PBPs are generally preferred because they: (a) incentivize performance rather than spending; (b) reduce the government's exposure to cost-accounting disputes; (c) create natural checkpoints that combine financial and programmatic oversight; and (d) can be structured to match contractor cash-flow needs, mitigating a key concern that contractors raise when asked to accept FFP. A well-designed PBP schedule should contain five to seven events per major performance period, with payment values that reflect the cost-incurrence profile of the work and do not exceed 90% of the contract or delivery-item price (FAR 32.1004).

### **Risk–Benefit Tradeoff Analysis**

Table 4 presents the risk–benefit tradeoffs that acquisition teams should weigh when evaluating CP-to-FFP conversion. The analysis is organized around seven dimensions of program impact.





**Table 4. Risk–Benefit Tradeoff Analysis for CP-to-FFP Conversion**

| Dimension               | Effect of FFP Conversion                                                                             | Risk Consideration                                                                  |
|-------------------------|------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Mission Resilience      | Positive: Incentivizes contractor investment in robust processes, supply chains, and quality systems | Risk: Contractor may cut corners to preserve margin if price is unrealistic         |
| Portfolio Flexibility   | Positive: Known price ceiling frees obligation authority for reallocation across portfolio           | Risk: Less flexibility to redirect effort mid-contract without costly modifications |
| Speed of Fielding       | Positive: Reduced administrative cycle time; fewer compliance gates                                  | Risk: Negligible if conversion is premature and triggers renegotiation              |
| Contractor Incentives   | Positive: Maximum profit motive to control costs and deliver efficiently                             | Risk: May incentivize narrow scope interpretation and aggressive ECP pricing        |
| Administrative Workload | Positive: Eliminates ICS, may remove EVM; reduces audit burden                                       | Risk: Shift to acceptance/quality oversight requires different skill sets           |
| Industrial-Base Health  | Positive: Rewards efficient performers with higher margins                                           | Risk: Unrealistic FFP can cause losses, workforce attrition, or default             |
| Rework/Default Risk     | Low if readiness gates are satisfied; contractor bears cost of rework                                | High if premature: cost overruns exceed margin, leading to performance degradation  |

## Oversight, Audit Burden, and Administrative Implications

One of the most frequently cited advantages of FFP conversion is the reduction in administrative burden. Under cost-reimbursement contracts, contractors must maintain approved cost-accounting systems subject to Cost Accounting Standards (CAS) compliance, submit annual incurred-cost proposals, and cooperate with DCAA audits that can take years to close. The government, in turn, dedicates contracting officers, cost analysts, and auditors to these processes. When a program converts to FFP, incurred-cost submissions are eliminated because the government’s payment obligation is fixed by the contract price rather than determined by actual costs (Barlow, 2026).

However, the shift from cost-type to fixed-price oversight is not a reduction in oversight—it is a transformation. Under FFP, the government’s focus must shift from cost tracking to technical performance, delivery quality, and acceptance protocols. The contracting officer must ensure that acceptance criteria are clearly defined in the contract, that quality-assurance surveillance plans are updated to reflect fixed-price metrics, and that the government retains sufficient technical access to evaluate deliverables. As Hastings et al. (2016) note, government oversight in a fixed-price environment requires learning to advise the contractor and understand potential risks, as opposed to performing the detailed cost-based oversight characteristic of cost-reimbursement contracts.

Several specific administrative adjustments are required during conversion. First, cost-type clauses (such as FAR 52.216-7, Allowable Cost and Payment) must be replaced with appropriate FFP clauses (FAR 52.232-1–52.232-5). Second, financing provisions must be realigned—either progress payments (FAR 52.232-16) or performance-based payments (FAR 52.232-28 and 52.232-32) should be incorporated, but not both. Third, EVM clauses should be removed unless the program’s risk profile warrants continued earned-value tracking—a decision that should be documented in the acquisition strategy. Fourth, the government should consider whether cost and technical data reporting—such as Contractor Cost Data Reporting (CCDR) for the DoW or Contractor Technical Data Reports (CTDR) for intelligence-community programs—should be retained for price-analysis purposes even under FFP. A clause-mapping sheet that



documents pre-conversion and post-conversion FAR/DFARS clauses is essential to prevent compliance gaps (Barlow, 2026).

Finally, the profit analysis must be adjusted. Under DFARS 215.404-4, the Weighted Guidelines method is used to develop prenegotiation profit objectives for DoW contracts requiring certified cost or pricing data. The contract-type risk factor within this methodology explicitly recognizes that FFP contracts warrant higher profit rates because the contractor assumes greater cost risk. This alignment between risk acceptance and profit opportunity is fundamental to the incentive logic of CP-to-FFP conversion—and should be communicated clearly to industry partners to encourage their acceptance of firm pricing.

## **Risks, Limits, and Counterarguments**

This paper has argued that disciplined CP-to-FFP conversion can yield significant benefits in cost discipline, administrative efficiency, and mission resilience. However, intellectual honesty requires acknowledging the substantial risks and limitations that attend any contract-type transition.

### **FFP Does Not Inherently Save Money**

The literature is clear that fixed-price contracts do not automatically reduce costs. Multiple authoritative studies—including the 2013 and 2014 Performance of the Defense Acquisition System reports, and the Arnold et al. (2008) IDA study—found no statistically significant difference in cost growth between fixed-price and cost-type contracts. The GAO reinforced this point in 2021, finding that although the DoW had substantially increased its use of fixed-price incentive contracts, it had not assessed whether this shift actually improved cost or schedule performance (GAO, 2021). MacDonald (2024) extends this insight, demonstrating that contract type is only one variable in a complex system where transaction costs, incentive misalignments, and compliance overhead collectively determine whether defense spending generates value. Fixed-price bids typically include higher risk margins (10%–12% or more compared to 5%–8% for cost-plus), and these margins represent a real cost to the government even if the contractor performs efficiently. The value proposition of FFP is not lower price per se; it is price certainty, incentive alignment, and reduced administrative burden.

### **Premature FFP Increases Risk and Cost**

When FFP is imposed before readiness conditions are met, the consequences can be severe. The GAO's analysis of Navy shipbuilding programs found a persistent pattern in which FPI contracts on lead ships with immature designs produced cost overruns exceeding ceiling prices, effectively shifting risk back to the government through follow-on contract adjustments and claims (GAO, 2017). MacDonald (2024) contextualizes these failures within a broader financial framework, noting that the defense sector faces unique bureaucratically imposed financial risks and an acquisition process that is too long and complex, offering poor returns for the risks taken. When fixed-price terms compound these existing structural pressures, the result is contractor financial distress, workforce attrition, and degraded government-contractor relationships. Congress recognized these dangers by enacting Section 808 of the FY2023 NDAA, which limits LRIP lots to no more than one when a fixed-price contract at Milestone B includes both development and production scope, with waivers requiring service acquisition executive approval and congressional notification (Pub. L. 117-263, § 808). The historical failure of Total Package Procurement in the 1960s provides a broader cautionary lesson about the consequences of forcing fixed-price terms onto immature programs.

### **Reduced Cost Visibility**

Under FFP, the government generally has no contractual right to access the contractor's actual cost data. This can create downstream challenges for pricing subsequent contracts,



understanding cost drivers for future program planning, and maintaining institutional knowledge of realistic cost estimates. MacDonald (2024) highlights the importance of cost transparency in defense budgeting, noting that the inability to accurately compare effective spending levels undermines informed decision-making. This principle applies equally at the contract level: without access to actual cost data, the government's ability to negotiate fair prices for follow-on efforts is significantly diminished. Acquisition teams should mitigate this by negotiating appropriate data-access provisions, retaining cost-reporting requirements (such as CCDR or CTDR) where justified, and ensuring that the pricing basis for the initial FFP conversion is well documented for use in future procurements.

### **Innovation and Development Should Often Remain Cost-Type**

This paper's framework explicitly recognizes that not all work should be converted to FFP. Innovation-heavy development, early-stage prototyping, efforts with undefined interfaces, and programs facing significant technology-maturation challenges should generally remain under cost-type contracts. FAR 16.301-2 authorizes cost-reimbursement contracts for these circumstances, and the Defense Business Board (2010) concluded that cost-reimbursement contracts are preferred for development efforts, particularly major weapons systems. Section 808 of the FY2023 NDAA codified this principle for major defense programs by limiting the number of LRIP lots on fixed-price contracts that combine development and production scope, unless the service acquisition executive grants a waiver (Pub. L. 117-263, § 808). The no-go pathway in this framework is not a failure—it is a disciplined acknowledgment that the wrong contract type can be more damaging than no change at all.

### **Scope and Competition Risks**

Converting a contract's type via modification raises scope and competition concerns. Material changes in type, scale, period, or cost may trigger out-of-scope findings that require recompetition or justification and approval under CICA. Even if the modification is legally supportable, the perception of a sole-source conversion can invite protest or congressional scrutiny. Thorough scope memoranda and legal reviews are essential safeguards against these risks.

### **Conclusion**

The question facing defense acquisition professionals is not whether to use firm-fixed-price contracts but when and how to transition to them responsibly. This paper has argued that CP-to-FFP conversion should be treated as a selective, readiness-based decision rather than a blanket policy preference. The eight readiness gates defined here—requirement maturity, configuration stability, interface stability, technical and integration risk, production and process repeatability, cost-data transparency, supply-chain robustness, and funding stability—provide a structured, evidence-based framework for making this determination.

The five conversion pathways—progressive risk transfer, CP-to-FPEPA, hybrid CLIN models, direct conversion, and the no-go pathway—offer practitioners a menu of options calibrated to different technical and market conditions. The practitioner toolkit, including checklists, CLIN patterns, financing comparisons, and risk–benefit tradeoffs, translates policy intent into executable acquisition strategy.

Properly executed, CP-to-FFP conversion does more than shift cost risk to industry. It frees obligation authority, reduces administrative overhead, sharpens contractor incentives, and enables the government to redirect scarce oversight talent from compliance processes toward mission-critical technical and quality evaluation. In an era when senior leaders are calling for the acquisition enterprise to move from bureaucratic risk aversion to a warfighting posture,



disciplined contract-type conversion is a practical tool for delivering “good enough now” capability rather than chasing perfect, slow, and operationally irrelevant solutions.

The framework advanced here is intentionally conservative: it identifies conditions under which continued cost-type contracting is the correct answer, not just conditions under which FFP is appropriate. This balance reflects the central insight of the literature—that contract type is not a panacea, and that the acquisition environment, incentive structure, and government-contractor relationship matter at least as much as the contract-type label. The goal is not to maximize the number of FFP contracts but to ensure that every contract-type decision is deliberate, evidence-based, and aligned with the overarching objective of delivering capability to the warfighter.

## References

- Arnold, S. A., Harmon, B. A., Tyson, K. W., Fasana, K. G., & Wait, C. S. (2008). *Defense department profit and contract finance policies and their effects on contract and contractor performance* (IDA Paper P-4284). Institute for Defense Analyses. <https://apps.dtic.mil/sti/pdfs/ADA482081.pdf>
- Barlow, M. J. (2026). *A readiness framework for converting contracts from CP to FFP: A Better Buyers module* [Presentation slides]. The Aerospace Corporation.
- Defense Business Board. (2010). *Best business practices for fixed-price contracting* (Report FY10-3). DoD. <https://dbb.defense.gov>
- Defense Federal Acquisition Regulation Supplement, 48 C.F.R. ch. 2 (2025). <https://www.acquisition.gov/dfars>
- Federal Acquisition Regulation, 48 C.F.R. ch. 1 (2025). <https://www.acquisition.gov/far>
- GAO. (2017). *Navy shipbuilding: Need to document rationale for the use of fixed-price incentive contracts and how they differ from cost-plus contracts* (GAO-17-211). <https://www.gao.gov/products/gao-17-211>
- GAO. (2020). *Cost-type contracts: Procedures needed for sharing information on contract choice among military departments* (GAO-20-352). <https://www.gao.gov/products/gao-20-352>
- GAO. (2021). *Fixed-price-incentive contracts: DOD has increased their use but should assess contributions to outcomes* (GAO-21-181). <https://www.gao.gov/products/gao-21-181>
- GAO. (2024). *Weapon systems annual assessment: Programs need to improve cost, schedule, and performance outcomes* (GAO-24-106831). <https://www.gao.gov/products/gao-24-106831>
- Hastings, S. E., Callaway, M. D., & Moeller, A. F. (2016). *Acquisition guidance for affordability: Using fixed-price contracts as a contributing tool for successful cost control* (TOR-2016-02779). The Aerospace Corporation.
- Office of the Under Secretary of Defense for Acquisition, Technology and Logistics. (2016). *Guidance on using incentive and other contract types*. DoD. [https://www.acq.osd.mil/dpap/cpic/cp/docs/Guidance\\_on\\_Using\\_Incentive\\_and\\_Other\\_Contract\\_Types.pdf](https://www.acq.osd.mil/dpap/cpic/cp/docs/Guidance_on_Using_Incentive_and_Other_Contract_Types.pdf)
- MacDonald, T. (2024). Using financial analysis to compare defense budgets. *Journal of Accounting and Finance*, 24(2), 110–126.



- Peters, H. M. (2024). *Department of Defense contract pricing* (CRS Report R47879). Congressional Research Service. <https://crsreports.congress.gov/product/pdf/R/R47879>
- Pub. L. 111-23, Weapon Systems Acquisition Reform Act of 2009, 123 Stat. 1704 (2009). <https://www.congress.gov/111/plaws/publ23/PLAW-111publ23.htm>
- Pub. L. 117-263, § 808, James M. Inhofe National Defense Authorization Act for Fiscal Year 2023, 136 Stat. 2395 (2022).
- Secretary of War. (2025, November 7). *Transforming the defense acquisition system into the warfighting acquisition system to accelerate fielding of urgently needed capabilities to our warriors* [Memorandum]. DoD. <https://media.defense.gov/2025/Nov/10/2003819439/-1/-1/1/>
- Section 809 Panel. (2019). *Report of the Advisory Panel on Streamlining and Codifying Acquisition Regulations* (Vols. 1–3). DoD. <https://discover.dtic.mil/section-809-panel/>
- Undersecretary of Defense for Acquisition, Technology, and Logistics. (2013). *Performance of the defense acquisition system: 2013 annual report*. DoD.
- Undersecretary of Defense for Acquisition, Technology, and Logistics. (2014). *Performance of the defense acquisition system: 2014 annual report*. DoD.
- Undersecretary of Defense for Acquisition, Technology, and Logistics. (2015). *Implementation directive for Better Buying Power 3.0*. DoD. [https://www.acq.osd.mil/fo/docs/betterBuyingPower3.0\(9Apr15\).pdf](https://www.acq.osd.mil/fo/docs/betterBuyingPower3.0(9Apr15).pdf)





## PANEL 5. HIGH-GROUND INNOVATION: SYNCHRONIZING AGILE DEVELOPMENT AND PRIVATE INVESTMENT IN SPACE

Wednesday, May 6, 2026, 1315 – 1430 ET (1015 - 1130 PT)

### Panel Summary:

Maintaining military preeminence in the space domain requires a fundamental shift in how the Department of War acquires and integrates orbital capabilities. This panel explores the dual pillars of modern space acquisition: the implementation of agile, DevSecOps, and digital engineering processes to streamline system development, and the strategic alignment of private-sector defense products with National Security Space needs. By examining case studies in AI-enhanced systems engineering and the efficacy of U.S. Space Force planning, these researchers offer a roadmap for delivering mission-critical space capabilities at the speed of relevance.

**Chair: Steven Butow**, Director, Space Portfolio, Defense Innovation Unit (DIU), Office of the Secretary of War (OSW)

**Discussant: Dr. Wenschel Deng Lan**, Chair, Space Systems Academic Group, Naval Postgraduate School

### Presenters:

**Accelerating DoW Space-based Acquisition through the Implementation of Agile, DevSecOps and Digital Engineering Processes** – *Michael Orosz, Research Director and Professor, University of Southern California Information Sciences Institute*

**Aligning Defense Solutions to National Security Space Needs** – *Andrew Berglund, Senior Analyst, Center for Space Policy and Strategy, The Aerospace Corporation*



**Steven Butow—General** Butow joined the team in 2015 as an early plank owner when DIU was known as the Defense Innovation Unit Experimental (DIUx). He initially served as West Coast Military Lead and helped establish the organization's operating locations in Silicon Valley, California; Boston, Massachusetts; and Austin, Texas. In 2016, he was selected as DIU's first Space Portfolio Director. Upon promotion to brigadier general, Butow continued this role at DIU as a DoD civilian.

General Butow continues to serve as the Commander of the California Air National Guard in a dual Title 10/32 status where he leads 4,900 Airmen assigned to five wings and a headquarters staff performing air, space and cyber operations as part of the Total Air Force. He commanded the 129th Rescue Wing at Moffett Field, CA, from 2011 to 2014.

General Butow's previous Title 10 leadership assignments include Deputy Director of the Joint Search and Rescue Center for U.S. Central Command in 2005, and as Chief of Personnel Recovery for U.S. Air Forces Central in 2007 supporting Operations IRAQI and ENDURING FREEDOM. Prior to joining DIU, Butow was the Vice Chief of the Joint Staff, California Military Department with responsibilities including cybersecurity, incident awareness and innovation. He has deployed ten times to support combat and contingency operations across Southwest Asia, the Middle East, the Horn of Africa and other global locations. General Butow has more than 3,500 flying hours in T-37, T-38, C-130, HC-130 and MC-130P aircraft.

Butow graduated from San Jose State University with a B.A. in Physics & Astronomy and earned a M.S. in Management with specialization in Air and Space Strategic Studies from the University of Maryland. He has also completed executive courses at Harvard's John F. Kennedy School of Government, and was a



member of the National Leadership Preparedness Initiative (NPLI) Cohort 16. Butow is a lifetime member of the Honor Society of Phi Kappa Phi for academic achievement.

As a researcher with the Search for Extraterrestrial Intelligence (SETI) Institute, Butow worked on instrument concepts for Mars surface soil analysis at NASA Ames Research Center in Mountain View, CA. He later served as co-principal investigator for a series of airborne science missions for which he was recognized for outstanding achievement and contributions to the Space Science Division in 1999 and received an Ames Honor Award as a member of an Astrobiology Mission Project Team in 2000.



**Dr. Wenschel Deng Lan**—is an Associate Research Professor at the Naval Postgraduate School in Monterey, California. A graduate of California Polytechnic State University, San Luis Obispo, she received Bachelor of Science and Master of Science degrees in aerospace engineering in 2008. At Cal Poly, she was on the team that launched the first CubeSats on both Russian and U.S. launch vehicles. Starting in 2007, she worked at the Rocket Systems Launch Program Engineering Center in San Bernardino as a mechanical engineer performing independent verification and validation for target and space vehicles. Lan returned to academia in 2010 to pursue a PhD in aerospace engineering at NPS. As a research associate in the Space Systems Academic Group, she continued research in enabling multi-manifest CubeSat launch opportunities. Recognized as a subject matter expert in aerospace structural analysis and dynamics, she has consulted

for various projects in the CubeSat and space launch communities. After earning her PhD in 2015, Lan continues to lead research of national interest and educate U.S. military officer students in space systems engineering and operations at NPS. Her teaching and research focus areas include spacecraft and payload design, augmenting national capabilities with small satellites and near-space platforms, and space applications for emerging technologies.



# **Accelerating DoW Space-Based Acquisition Through the Implementation of Agile, DevSecOps and Digital Engineering Processes**

**Michael Orosz**— directs the Decision Systems Group at the University of Southern California's Information Sciences Institute and is a Research Associate Professor in USC's Sonny Astani Department of Civil and Environmental Engineering. Dr. Orosz has over 30 years' experience in government and commercial software development, systems engineering and acquisition, applied research and development, and project management, and has developed several successful products in both the government and commercial sectors. Dr. Orosz received his BS in engineering from the Colorado School of Mines, an MS in computer science from the University of Colorado, and a PhD in computer science from UCLA. [mdorosz@isi.edu]

**Lt Col John "Mack" Turner**—joined the U.S. Space Force in 2022 after 15 years in the U.S. Army. He has held all Artillery positions at company level to include both Battery and Company Command. After supporting PEO Ammunition, he joined PEO C3T modernizing Mission Command software. In his last assignment, he oversaw Artillery software requirements for the Fires Center of Excellence. Since entering the Space Force, Mack has supported the follow-on effort for the Next Generation Operational Control System (OCX 3F). He holds a BS in electronic and information technology systems and an MA in defense and strategic studies. [john.turner.66@spaceforce.mil]

**Lt Col Trevor Johnson**—is the Deputy Commander for the 31st Capability Development Squadron, Los Angeles AFB, California. He leads 35 military and civilian members in the development of the Next Generational Operational Control System. Lt Col Johnson earned his bachelor's degree in management from the U.S. Air Force Academy and was commissioned in 2006. He distinguished himself in a variety of leadership positions within the aircraft maintenance and acquisition career fields while at Grand Forks and Los Angeles AFB. Prior to assuming his current position, Lt Col Johnson led execution of the Future Operationally Resilient Ground Evolution (FORGE) program.

**Brian Duffy**—is a Senior Systems Engineer with the University of Southern California Information Sciences Institute (USC/ISI). He conducts research and analysis to determine system engineering methods and metrics necessary to transition Major Defense Acquisition Programs from a traditional waterfall development to Agile/DevSecOps processes. Prior to USC/ISI, Mr. Duffy retired from the U.S. Air Force with multiple assignments related to National Security Space acquisition programs and C2 systems. Mr. Duffy holds a Masters of Aeronautical Science from Embry-Riddle Aeronautical University and a Bachelors of Aeronautical and Astronautical Engineering from the University of Washington.

**Craig Charlton**—is a Senior Systems Engineer at USC's Information Sciences Institute (USC-ISI) and has provided acquisition support at Space Systems Command (SSC) at the Los Angeles Air Force Base during the past 20 years on several leading-edge satellite systems. Prior to his position at SSC, Mr. Charlton acquired more than 25 years of experience as a software engineer and in managing software projects in the commercial world, primarily in the fields of engineering and of law enforcement. Mr. Charlton received a BA in mathematics from California State University, Long Beach.

## **Abstract**

Over the past nine years, using funding from the U.S. Air Force and the U.S. Space Force, System Engineering Research Center (SERC) researchers at the University of Southern California's Information Sciences Institute (USC/ISI) along with the active support and engagement of the USAF and USSF have undertaken a series of case studies focused on accelerating the Department of War's space systems acquisition process through the application of agile, DevSecOps, digital engineering and other advanced technologies such as AI.

The major focus of this research is discovering challenges and exploring solutions to managing process flow, resources and deliverables throughout the agile/DevSecOps system development process. Such challenges include managing dependencies on external and internal systems; supporting real-time continuous multi-system integration and testing; managing staff loading and



specialties over the course of the program; adapting to new and changing capabilities; arranging operator engagement throughout the development process; and the availability, collection and analysis of performance metrics for improved situational awareness and real-time decision-making.

Several observations and lessons learned will be covered in this paper. In addition, this paper discusses lessons learned to date in employing AI as a systems engineering aid throughout the development and deployment cycle.

## Introduction

Secretary of War Pete Hegseth mandated that the U.S. Department of War (DoW; 2025) needs to field systems to the warfighter sooner, on budget and that meet warfighter needs. The purpose of the research undertaken by USC and the SERC UARC (Systems Engineering Research Center, n.d.) is to address this mandate by determining how agile (*Manifesto for Agile Software Development*, n.d.), DevSecOps (Microsoft Security, n.d.) and other modern acquisition methods and approaches—many already being successfully used in the private sector—can be successfully implemented in current and future DoW acquisition programs. Specifically, what tools and approaches can be incorporated into the acquisition process that allows programs to become agile and nimble to address dynamic conditions from both an operations perspective (i.e., what does the warfighter need now) and also from an acquisition perspective (e.g., how has technology changed, have new personnel joined the team, etc.)? In other words, by implementing processes which reduce program risk, the DoW will shift the discovery and resolution of program risk earlier in the program, leading to a more rapid delivery of warfighter capabilities.

The research discussed in this paper is currently focused on improving the acquisition of space-based software command and control (C2) systems, but much of what has been observed and learned can also be applied to hardware-only and hybrid hardware/software programs in other domains within the DoW, for example, agile processes applied to hardware-based programs that employ additive manufacturing (i.e., 3D-printing).

Over an approximate nine-year period, the USC-led project team has observed that agile, DevSecOps and related modern systems acquisition methods are able to improve DoW acquisition, though it is apparent that many challenges still exist and that adopting industry best practices requires considerable tailoring. For example, “code fast, fail fast and learn” is difficult to achieve in large mission-critical enterprise systems that cannot fail.

This paper extends lessons learned and observed from previously published work (Orosz, Hempen, et al., 2025; Orosz, Turner, et al., 2025).

## Process Used

As noted in Orosz, Hempen, et al., (2025), the USC-led team is embedded into the space-based C2 acquisition environment. Project team members serve two roles: first, as systems engineering subject matter experts (SMEs) on the government team (e.g., members of one or more Integrated Project Teams [IPTs]), and secondly, as researchers. The advantage of embedding into the day-to-day acquisition process is that the team is able to observe which agile and other modern systems development processes and approaches either work or fail in improving and accelerating systems development and delivery. This embedding has helped establish strong working relationships with both the government and prime contractor teams (responsible for the actual development of the system) which has allowed experimentation with different approaches (e.g., making recommendations to both government and prime contractor) in the quest to improve and accelerate the acquisition process. Figure 1 summarizes the various projects the team has supported as part of this research effort. Currently, the project team is



supporting Project D, and many of the research results presented in this paper are from this project. The team is not supporting Project E, but the project is listed here as it directly impacts Project D, as discussed later in the paper.

| Project   | Development Method           | Duration    | Lines of Code | Notes                                              |
|-----------|------------------------------|-------------|---------------|----------------------------------------------------|
| Project A | Waterfall                    | 39 Months   | 178K          | Serves as baseline for research                    |
| Project B | 50-50 Hybrid Agile/Waterfall | 25 Months   | 113K          | Agile effort greatly outperformed waterfall effort |
| Project C | Tech Study                   | 15 Months   | None          | Prep for Project D. Agile training                 |
| Project D | 70/30 Hybrid Agile/Waterfall | 75 Months+  | 150K+         | Currently underway                                 |
| Project E | Waterfall                    | Many months | N/A           | Impacts Project D                                  |

**Figure 1. Table Summarizing the Space-Based Command and Control Projects That the USC-Led Research Team Was or Is Currently Supporting. The research team is not involved with Project E; however, Project E is linked to Project D, as discussed in the paper.**

## Observations

Project D relies on a modified implementation of the SAFe® (Lindsey, 2026) agile approach to development. As noted in Orosz, Turner, et al. (2025), requirements are decomposed into one or more work packages called features which are to be accomplished within a 12-week period of development work within a program increment (PI). Features are then further decomposed into one or more work packages called stories which will accomplish the work required to complete the overall feature (Figure 2). Each feature is assigned a priority on the project backlog which coincides with the PI the feature is to be worked. Each program increment is 13 weeks in length and consists of four 3-week sprints followed by a 1-week non-development effort used for demonstrations, training, PI planning, retrospection and other activities (Figure 3).




**Figure 2. Project Consists of One or More System Requirements. Each requirement is decomposed into one or more features. A feature represents development work that can be completed within 12 weeks. Each feature is decomposed into one or more stories. Each story represents work that can be completed within weeks. Figure adapted from Orosz, Hempen, et al. (2025).**






**Figure 3. Features Are Assigned to a 13-Week PI. Within the PI are four 3-week sprints in which stories decomposed from a feature are worked. The final 1 week for each PI is used for non-development work such as training, innovation/research, retrospectives and other activities. Figure adapted from Orosz, Hempen, et al. (2025).**

Figure 4 shows the number of features planned per PI (black trace), the number of features that were completed (green trace) and the number of uncompleted features that slipped from that PI to future PIs (blue bars). Slipped features indicate situations where the development team was not able to complete their assigned feature within the assigned PI.




**Figure 4. Features Assigned, Completed and Uncompleted per 13-Week PI. Blue bars represent features that did not complete within the 13-week PI and “spilled” over into future PIs.**




**Figure 5. Plot of Features Completed per PI, as of October 2024. Note large number of features (a “bow wave” assigned to PI-15 [January 2025]). These features include uncompleted features that slipped into future PIs. All features are to be completed by the end of PI-17 (July 2025). R1-R5 are MVP/MMP releases milestones.**



**Figure 6. Upper Graph Represents Feature Work-Off Performance Approximately One Year Later Than That Covered in Figure 5. The development schedule has been extended to another three PIs (approximately 9 months) to PI-20. Milestone releases R4-R5 have shifted “right.” Bottom Graph Shows Bugs That Have Been Detected via Integration and Testing Efforts. Large upticks in discovered bugs come at the expense of feature work-off as key workforce members are pulled from development to address critical bugs (CAT1Us), shown in red.**

Figure 5 is a snapshot of program progress as of October 2024. The figure shows the number of features completed (dark blue bars) and the number of features planned for the current and future PIs (gray bars). As shown in the graph, the number of features assigned to be worked in the upcoming PI-15 (approximately 75), centered in January 2025, is considerably higher than feature velocity tracked from PI-0 through PI-14 (July 2021 to October 2024). This large number of features represents an evolving “bow wave” of work that has accumulated due to features slipping from PI to PI throughout the program.

At the top of Figure 6 is a similar plot as in Figure 5 but represents project performance approximately a year later, as of November 2025. The forecasted “bow wave” of features, centered around January 2025, has been “smoothed” out and shifted by extending the development schedule another three PIs (approximately 9 months) from PI-17 to PI-20. At the bottom of Figure 6 is the trace of “bugs” (known as CAT1s and CAT2s for “Category 1” [critical, with no work around] and “Category 2” [important, but not critical]) per PI. Notice that feature velocity is directly linked to the number of “bugs” detected and worked. The more bugs detected and assigned to be worked (i.e., CAT1U bugs), the lower the feature velocity.

The large number of bugs detected late in the development effort (starting roughly in PI-11) is attributed to the delayed start of integration and testing (to be discussed later). The CAT1U (critical with no work around) bugs have to be addressed, requiring key development team members to be pulled from feature development to work those bugs which is contributing to the low feature velocity that has been observed.

## Discussion and Recommendations

### What Is Working

Although the development schedule has slipped (at the time of this paper, approximately 9 months), there have been a number of “wins” observed (to date) in Project D. First, a major release was developed and underwent early testing by the end-user community via a near operational environment. The availability of this release (an MMP) was delayed by about 6 months due to challenges to be discussed below. Although the operating environment was simulated, this early testing and simulated use of the capability provided early feedback to the program which provided valuable input in the development of the next major MVP and MMPs.

A second win was the prime contractor’s management of the project backlog. When blocked features or stories were encountered, the development teams quickly pivoted and “pulled forward” and worked the next highest priority task from the project backlog. As blocked features and stories became unblocked, work once again resumed on those tasks based on their current priorities on the project backlog (which might have changed while blocked). Projects with a definitive end state (e.g., DD250, etc.), at some point, may be challenged to “pull forward” future high priority tasks as fewer tasks are available as the project backlog shrinks. Nonetheless, at the time of this writing, the development effort for Project D only slipped 9 months (i.e., 3 program increments) from the original 52-month (17 program increments) development schedule.

The third win was the unfettered access the government had to project technical data. The government had direct access to the project backlog, requirements and test databases throughout the program. In addition, the government teams participated in all scrums, ceremonies, demonstrations, and planning activities throughout the project. This unfettered access helped the government to have a real-time/near real-time picture of the current state of the project. Much of the data extracted from this access was formatted and analyzed by the government with tools developed using Python, MS Visual Basic for Applications, MS Excel and other applications. These tools allowed the government team to quickly analyze and provide



real-time performance metrics to senior government leadership. In many cases, this information was shared with the prime contractor to help identify challenges in the acquisition process and establish priorities to address. Figures 4, 5, and 6 from above are examples of the metrics generated to provide real-time status of the project.

Perhaps the biggest “what worked” to date in Project D was the ease in which the prime contractor’s scrum teams worked with the government’s IPTs. Through frequent engagements via scrum meetings, ceremonies and other activities, team members from both sides developed a relationship that promoted good communications and established a common teaming environment while still respecting the business needs of the respective teams (e.g., the government is a representative of the warfighter/end-user, the contractor is the hired vendor, etc.).

## Challenges

**Late Integration and Testing:** As noted in Orosz, Hempen, et al. (2025) agile-based projects need access to a near operational environment as soon as possible when implementing agile, DevSecOps and Continuous Integration/Continuous Deployment (CI/CD) approaches. The integration and testing in Project D were delayed due to the lack of availability of a near operational test environment. Of the testing that was undertaken early in the program, much of that testing consisted of isolated functional testing at the subsystem or Configuration Item (CI) level. Horizontal testing where one or more CIs are tested as an integrated unit (i.e., end-to-end testing) did not ramp up until around PI-11 (Figures 5 and 6 above). The result was a bow wave of “bug” (i.e., the CAT1Us and CAT2Us) discovery which led to a shifting bow wave of features to be completed (see Figures 5 and 6). In this case, the development team focused on the high-priority bugs rather than on the features and stories residing on the project backlog.

**Workforce Management in a Dynamic Environment:** A major challenge for the prime contractor is managing the workforce in an agile program. As program priorities change, often due to changing requirements and end-user needs, there is pressure (driven by these changing priorities) on the prime contractor to quickly pivot and focus on delivering newly identified high-priority capabilities. Often, however, the skill sets needed to meet these newly identified high priority capabilities may not be present in the current prime contractor’s workforce. The prime contractor may often be in the difficult position of having to work with the workforce which they currently have on hand. Their skill sets may be more aligned with previous high-priority and scheduled work which is now further down in the project backlog (i.e., scheduled to be worked in the future) due to the recent focus on the newer high-priority work.

In such circumstances, development progress (e.g., velocity, etc.) can dramatically slow as the existing workforce has to “ramp-up” and become more familiar with the newly classified high-priority work to be completed. Further, under these conditions, there is a tendency of the prime to pull forward (i.e., raise higher in the project backlog) lower priority work to be worked with the current available workforce (with the more aligned skillsets) at the expense of working on higher priority work that needs to be delivered early. In other words, the prime contractor does not want a highly skilled workforce working tasks that actually require lesser skills. This is another example of “assuming the workforce has equal skill levels and experience and can be used on any effort” that has previously been identified (Orosz, Turner, et al., 2025).

**Corporate Knowledge is Lost:** A key factor contributing to schedule slippage is the constant churn of the government and prime contractor workforce. For many DoW programs, including agile programs, the program can span multiple years, even decades. In fact, if DD250s (AcqNotes, 2024) and other key milestones are viewed as MMPs, programs can be viewed as continuous evolving enterprises. Once a DD250 has been reached, the next phase (often an upgrade or sustainment) starts. As the program progresses, prime contractors change,





government agencies restructure, developers come and go, and military and government civilians transition to other assignments. The challenge here is that knowledge of the system is often lost due to this churn in the workforce.

In this context, knowledge encapsulates the system being developed and sustained (i.e., system design, architecture, etc.), the acquisition process used to develop the system (i.e., waterfall, agile, hybrid, integration & testing approach, etc.) and the operational environment in which the system will operate (i.e., domain knowledge). Decisions behind design considerations, priorities, and other programmatic decisions are often lost when team members transition to other programs or opportunities. Often, priorities, stories, features, and tasks on program backlogs are set, and the rationale behind these decisions are lost as the program progresses. Capturing this corporate knowledge and making it available to new members of the workforce is a major challenge. Employing modern engineering processes can help (e.g., MBSE tools, etc.); however, often these tools and processes are not seamlessly interconnected. For example, linking system requirements to the model design to the developed software and system test results is still very difficult. In addition, capturing the decisions behind the design and approaches used is rarely captured, at least digitally, which can be problematic years into the future. Note that this “linking” challenge not only applies to the program office, but it also applies at the enterprise level.

**Implementing DevSecOps:** As noted in Orosz, Hempen, et al. (2025), in space-based enterprises, the implementation of DevSecOps and related Continuous Integration/Continuous Deployment (CI/CD) is challenging. This is also true in Project D. Often, the command and control (C2) segment of a space-based program is initiated and often completed prior to the availability of other systems and subsystems—including the space vehicle (SV) constellation. The lack of these external systems limits how much operational use (or testing) can be undertaken in the “Ops” portion of DevSecOps. This challenge also appears when external systems or their simulators are constantly changing or when commands or tasking (e.g., for commanding a space vehicle) frequently change. Enhanced enterprise engagement can help mitigate these challenges.

As noted elsewhere in the paper, project D suffered from the delayed availability of a near operational test environment that can be used for testing prior to deployment. In addition, the lack of available operators to test the incremental releases was also challenging to the project as these individuals are part of a small, highly trained cadre of operators that are in high demand. Scheduling time for these operators to test releases of the system was problematic and often involved “luck.”

**Loss of Architecture Modularity:** Despite an initial focus and the best intentions of the prime contractor throughout the program to develop and maintain a modular architecture, there were many examples of tight coupling between multiple modules and Configuration Items (CIs) within Project D. This does not imply that the prime contractor never undertook refactoring during the development of the system. Refactoring was, in many situations, necessary as new capability was added to the system. However, such refactoring was chosen as the path of least resistance when adding new capability to the system. As the program progressed, however, multiple engineering change orders were introduced (via Requests for Change [RFCs]), high priority bugs (or discrepancy reports [DRs]) were worked, new team members transitioned onto the project, senior systems engineering leadership changed, and the constant pressure of maintaining high development velocity and production of releases (MVPs, MMPs, etc.) resulted in the tight coupling between various modules (within and between CIs). Refactoring to maintain modular architecture and code baseline was important, but a lower priority.



**AI and the Acquisition Process:** The government and USC/ISI team undertook an initial exploration on using available Artificial Intelligence (AI) tools to help with program management and metrics tracking. The motivation behind this exploration was to reduce the workforce effort required to collect, aggregate and report on program performance metrics. As with many other programs, without an AI or some other automated or semi-automated tool, collecting, aggregating, analyzing and then reporting on various aspects of a program is predominantly a manual process which is time-consuming and often error prone. In Project D, the government explored using NIPRGPT (U.S. Department of the Air Force, 2024)—a Large Language Model (LLM)—to collect, aggregate, analyze and generate program progress reports. The reports produced by the LLM required some final editing by a human subject matter expert (SME), but the savings in workforce hours (100s of hours) was deemed a success.

## Recommendations and Observations

As previously noted, this paper summarizes recent research and findings to an ongoing series of case studies focused on accelerating an agile-based software program. Figure 7 summarizes recommendations and observations previously reported (Orosz, Turner, et al., 2025). The recommendations that follow are an extension to these earlier observations.

| Recommendation/<br>Observations    | Comments                                                                                                                                                                                                                                                                                                                       |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Upfront engineering                | Perform some upfront engineering to help populate the project backlog, map features with compliance requirements, identify test cases, and identify dependencies as early as possible in the program                                                                                                                           |
| Not all staff skillsets are equal  | The amount and diversity of talent needs to be constant, so they can be moved/re-arranged to react to the changing needs (be agile) of the program                                                                                                                                                                             |
| Near operational test environments | Establish (early in the program) a near operational environment and high-fidelity simulators (for horizontal I&T)                                                                                                                                                                                                              |
| Sprint margins                     | Plan margin into the sprints to handle unexpected events such as new technology insertion and/or unexpectedly complex stories (Agile 101)                                                                                                                                                                                      |
| Licensing and other IP issues      | Get licensing, IP, accreditation, certification and other programmatic issues resolved early                                                                                                                                                                                                                                   |
| Training                           | Need for on-boarding and continuous training to ensure team members (both the contractor and acquisition team) are on the same page                                                                                                                                                                                            |
| Costing challenges                 | For projects that are costed using “T-shirt sizes,” costing needs to be segmented into development, integration and testing, and bug (problem) work-off. Also, need to have a sufficient number of T-shirt sizes to cover the smallest to largest work packages                                                                |
| Custom tools                       | Be prepared to customize performance tracking tools – Applies to all teams – Issues: <ul style="list-style-type: none"> <li>• Software incompatibilities</li> <li>• Foreign ownership of tools</li> <li>• Access challenges (e.g., VPN, security concerns, etc.)</li> <li>• EVM lags actuals – sometimes by months.</li> </ul> |

**Figure 7. Summary of Lessons Learned and Observations From Previous Research (Orosz, Turner, et al., 2025)**



**Capturing Corporate Knowledge:** Recognizing that workforce churn will probably not abate even with increased incentivization as career growth and other opportunities will continue to drive employees to transition off a project, it is recommended that there should be a focus on 1) improving the capture of corporate knowledge that can be shared and 2) providing relevant on-boarding and continuous workforce training. The following are some recommendations for achieving this goal:

1. As features and stories are added, worked, and completed on the project backlog, design and coding decisions (from scrums, demonstrations, ceremonies, retrospectives and other meetings) should be captured and stored with the tracking ticket (e.g., a Jira [Atlassian, n.d.] Ticket) on the backlog. Such information could include a few bullets of text or, in environments that can support voice-to-text technologies (e.g., non-classified environments), a textual record of these meetings. For example, in programs that can support VTC meetings using tools such as MS Teams (“Microsoft Teams,” 2026), the tool itself can capture voice traffic and automatically translate a textual record that can be attached to a tracking ticket. The advantage of having access to this information is that the thought process behind various decisions in adding, implementing, testing and then closing a feature or story is available to the future workforce.
2. Maintain and frequently update on-boarding training materials that describe the development approach used (e.g., SAFe, scrum, Kanban, etc.), design decisions behind the system, and other programmatic information that the workforce member needs to know. In addition, this on-boarding and continuous-training material should include domain-specific information to help the development and government team members to better understand the targeted end-user environment. For example, in a space-based command and control acquisition program, training (on-boarding and frequent refreshers) that highlights space-to-ground communications, cyber vulnerabilities, space vehicle constellation dynamics, orbital dynamics and other related information can help focus developers and the government on end-user needs and priorities.
3. Increase domain knowledge. A possible pathway is to rely on a qualification system similar to military aviation and other mission-critical job functions. For example, with aviation, pilots must not only achieve certain qualifications, but they also must maintain them via frequent “refresher” flights or training. The same could be used in the acquisition environment with domain-specific qualifications (e.g., space-based launch operations) and acquisition-specific qualifications (e.g., current agile best practices, etc.).

**Frequent Engagement With the Enterprise:** Team members (developers, integrators, testers, chief engineer, senior leadership and others) from both the prime contractor and the government need to be in frequent engagement with the whole enterprise (space segment, control segment, user segment, etc.). Whether this is in the form of working groups, tech exchange meetings (TEMs) or as part of other frequently scheduled meetings (e.g., project management reviews [PMRs]), they need to occur frequently to allow all stakeholders an opportunity to become aware of possible changes coming in the future, the status of the development/testing environments, and the overall status of the enterprise. In short, changes to the system design, requirements or concept of operation should not be a surprise to the development and testing environment.

Successful engagement with the enterprise is a particularly challenging problem with prime contractors due to a lack of mission expertise and historical links to the government. It is a surprisingly difficult task to engage with the enterprise. It is much easier to be successful at this task if the prime contractor has hired individuals who have prior relationships and are trusted as mission experts. Without this relationship, contractors are seen as “outsiders” who “don’t



understand the mission” and enterprise engagements will not be fruitful. Additionally, it is hard to do the upfront engineering of the mission threads without this expertise, and it is much harder to do integration without knowledge of where the “gotcha’s” will come and without relationships with the other entities that need to be integrated with. The best software experts in the world will fail if they do not understand the mission.

**Trained and Experienced Workforce.** A trained and experienced project workforce is critical to accelerating an agile program. As noted in Orosz, Turner, et al. (2025), often both the government and the prime contractor view the workforce as a “plug and play” commodity where all team members have roughly the same experience and skill levels. This is simply not the case nor is the issue simply focused on the level of software development experience. Knowledge of the domain (space operations, ConOps, etc.), knowledge of development and testing environments, and awareness of best practices are also essential in the workforce.

As noted, the workforce is not a commodity. It is very difficult to hire a replacement that can be “plugged” into the development environment. Even experienced personnel will require time to “learn” the development environment, the ConOps of the targeted application, and workflow of the integration, testing and deployment. In project D, it was common for the prime contractor and the government to take up to 12 months to on-board new (and often) experienced personnel. It is recommended that both on-boarding and continuous training be an integrated part of the acquisition workflow. Although an expense, the investment can be spread over the life cycle of the program which often spans multiple years and, in some cases, decades (e.g., space-based systems).

**Integration and Testing:** As noted in Orosz, Turner, et al. (2025) and summarized in Figure 7, the sooner a near-operational environment becomes available, the sooner system-wide integration and testing can be undertaken. Equally important, however, is the need to define up front use cases or “mission threads” that define how the system works (or is envisioned to work) and is used. These cases included expected inputs, output, error and exception handling are directly linked to system requirements. These use cases inform the development of test scripts that are used during system-wide integration and testing. Additionally, including use case testing in a government statement of work could enable integrated “graduation” testing for releases, especially when a functional qualification test per release is not feasible.

Although the above recommendation is not new, it is often overlooked in agile-focused programs where the “code fast, fail fast and learn” approach is expected to drive the system design and test environment. This simply is not the case in complex OSW programs such as command and control of a space vehicle constellation where there are multiple vendors and development approaches and timelines that span the enterprise.

For example, in many of these OSW agile-based software development programs, the development prime contractor will assume that the successful testing and verification of all features/stories (i.e., work products) that are decomposed from a single system requirement indicates that the system requirement has been met. Unfortunately, the sum of the parts does not necessarily equal the complete system. Only when all components are tested as an enterprise system, where all modules and CIs are connected and operational (even if simulated) and subjected to successful end-to-end testing, can it be declared that the system requirement(s) have been satisfied.

**Increased Government Involvement:** One of the successful aspects of all the projects that the research team has been involved with is the value of embedding knowledgeable government acquisition team members within the development environment. Government team members participate in scrums, ceremonies, planning, testing, and demonstration events. Such



activities help build trust between the prime contractor and the government. For example, the government is aware and can appreciate challenges faced by the prime contractor (e.g., workforce changes, budget and scheduling challenges, code complexity issues, etc.). Similarly, the prime has an opportunity to appreciate the challenges faced by the government (especially the warfighter). Issues such as the importance of meeting system requirements, what those system requirements mean to the warfighter and the urgency of getting capability into the hands of the warfighter help with managing the project backlog, workforce and other resources.

In addition, this embedding of the government into the development workforce often allows the government real-time or near real-time access to prime contractor data that is used for tracking/capturing project performance metrics.

That said, government involvement requires a knowledgeable government workforce. As previously noted, a major challenge in agile-based project efforts is the need for an experienced workforce. Throwing a freshly minted lieutenant with no prior experience or knowledge domain into a scrum will almost always hinder system development.

**Maintaining MOSA:** As noted previously in the paper, the prime contractors behind each of the projects—including Project D—did focus on developing and maintaining a modular open system architecture (MOSA). The challenge is in maintaining modularity as the program progresses. Often, as pressure to meet upcoming milestones intensifies, the focus shifts to getting the capability completed, regardless of how it is done. This approach often results in unintentional interdependencies within and between modules. Although often discussed in the systems and software engineering community (Fowler, 2002; Schults, 2022), refactoring is often a second thought as such efforts add costs and time to a project. Although Fowler (2002) and Schults (2022) do not recommend adding refactoring stories or tasks specifically to the project backlog due to the challenges noted above, particularly when addressing systems architecture, it is recommended that refactoring stories/tasks be included in the project backlog. If not included in the project backlog, then at a minimum, the contracting language should require the prime contractor to prove that modularity is being maintained in the project. In other words, refactoring is just as important as the highest priority stories/tasks on the project backlog.

## Conclusion

The purpose of the research undertaken by USC and the SERC UARC is to determine how agile, DevSecOps and other modern acquisition methods and approaches can be successfully implemented in current and future DoW acquisition programs to accelerate systems development and deployment. Specifically, the research is focused on determining what tools and approaches can be incorporated into the acquisition process that allows programs to become sufficiently agile and nimble to address dynamic conditions from both an operations perspective (i.e., what does the warfighter need now) and also from an acquisition perspective (e.g., what tech has changed, has new personnel joined the team, etc.).

A series of recommendations and observations have been produced from this and past studies. Chief among those recommendations is the need to do some up-front systems engineering to ensure that the necessary “scaffolding” is in place prior to the start of an agile program. In addition, it is important that a near-operational integration and testing environment be put in place as soon as possible to allow DevSecOps and CI/CD operations to be implemented. Another recommendation is that a trained workforce is absolutely necessary for the successful implementation of an agile-based program. This requires having the on-boarding and continuous training materials on-hand that cover programmatic details of the acquisition environment and the targeted domain.





As for next steps, the team will continue to explore introducing and implementing agile, DevSecOps and other modern systems development approaches into the DoW acquisition environment. Of particular interest is the challenge of transitioning from a completely virtual model of the proposed system to a final deliverable system by replacing the virtual components with the actual working components of the final system. An additional active area of research by the USC/ISI team is in employing AI in the capturing and dissemination of corporate knowledge. An open question is whether an LLM or other AI technique, when exposed to various system artifacts, can provide the necessary information when used to on-board and train new and existing workforce members.

## References

- AcqNotes. (2024, February 23). *Developing DD form 250*. <https://acqnotes.com/acqnote/careerfields/departments-of-defense-form-250>
- Atlassian. (n.d.). *Jira: Issue & project tracking software*. <https://www.atlassian.com/software/jira>
- Fowler, M. (2002). Refactoring: Improving the design of existing code. In *Lecture notes in computer science* (p. 256). [https://doi.org/10.1007/3-540-45672-4\\_31](https://doi.org/10.1007/3-540-45672-4_31)
- Lindsey. (2026, January 21). *Framework - Scaled Agile framework*. <https://framework.scaledagile.com/#big-picture>
- Manifesto for Agile software development*. (n.d.). <https://agilemanifesto.org/>
- Microsoft Security. (n.d.). *What is DevSecOps? Definition and best practices*. <https://www.microsoft.com/en-us/security/business/security-101/what-is-devsecops#:~:text=DevSecOps%2C%20which%20stands%20for%20development,of%20a%20multicloud%20security%20strategy>
- Microsoft Teams. (2026, January 24). In *Wikipedia*. [https://en.wikipedia.org/wiki/Microsoft\\_Teams](https://en.wikipedia.org/wiki/Microsoft_Teams)
- Orosz, M. D., Hempen, J., Turner, J. M., Duffy, B., & Charlton, C. (2025). *Lessons learned in implementing DOD space-based Agile/DevSecOps programs*. <https://gsaw.org/wp-content/uploads/2025/10/2025-S12-2-Lessons-Learned-in-Implementing-DoD-Space-based-Agile-DevSecOps-Programs.pdf>
- Orosz, M., Turner, M. J., Pool, J., Duffy, B., & Charlton, C. (2025). *Lessons learned in implementing integration and testing in a DoD space-based Agile/DevSecOps environment*. [https://www.dropbox.com/scl/fo/bwqt81qgp7gkznr0e38i6/AM1xi0q5QpOqPjP488FvoS0/Papers%20-%2034th%20Aerospace%20Testing%20Seminar?dl=0&preview=Michael+Orosz\\_Agile+Aerospace+Testing+Seminar\\_Paper\\_20250313+\(1\).pdf&rlkey=t80lb3igo9bgptii1wv7vu935&subfolder\\_nav\\_tracking=1](https://www.dropbox.com/scl/fo/bwqt81qgp7gkznr0e38i6/AM1xi0q5QpOqPjP488FvoS0/Papers%20-%2034th%20Aerospace%20Testing%20Seminar?dl=0&preview=Michael+Orosz_Agile+Aerospace+Testing+Seminar_Paper_20250313+(1).pdf&rlkey=t80lb3igo9bgptii1wv7vu935&subfolder_nav_tracking=1)
- Schults, C. (2022, January 10). *The when, why, and how of refactoring in Agile*. <https://linearb.io/blog/refactoring-in-agile>
- Systems Engineering Research Center. (n.d.). <https://sercuarc.org/>
- U.S. Department of the Air Force. (2024, June 10). *Department of the Air Force launches NIPRGPT*. <https://www.af.mil/News/Article-Display/Article/3800809/departments-of-the-air-force-launches-niprgpt/>
- U.S. Department of War. (2025, November 10). *Secretary of War announces acquisition reform*. <https://www.war.gov/News/Releases/Release/Article/4329487/secretary-of-war-announces-acquisition-reform/>



## Aligning Defense Products to National Security Space Needs

**Andrew Berglund**—is a senior policy analyst within the national security team at the Center for Space Policy and Strategy at The Aerospace Corporation. In his role, he covers issues related to national security space, defense acquisition and budget policy, and emerging technologies. Prior to joining Aerospace, he was a senior defense acquisitions analyst at the U.S. Government Accountability Office, where he earned a meritorious services award and led work on national security space, acquisition policy, cybersecurity, and nuclear command and control. His prior federal service also includes positions at the U.S.-China Economic and Security Review Commission and the Office of Naval Intelligence. He earned a bachelor's degree in political science from Christopher Newport University and a master's degree in global policy studies from the LBJ School of Public Affairs at the University of Texas at Austin. [andrew.berglund@aero.org]

### Abstract

The U.S. defense acquisition system has struggled to keep pace with security threats and rapid technology change, despite decades of reform efforts to integrate commercial innovation in military systems. Traditional acquisition approaches with long development timelines tied to predetermined and static system requirements do not meet the urgency or speed that military leaders and political decision-makers demand. Powered by software-driven business models and private capital, a cohort of defense tech companies have developed new approaches to meet both military and commercial needs with dual-use products. However, military users will always require a set of capabilities that differ in important and substantive ways from what the commercial market is able to deliver. This gap has contributed to a small but important new segment of the defense industrial base, companies that self-fund development of military-specific systems and capabilities in anticipation of government demand and funding. If the U.S. Space Force correctly signals its future demand and adapts its processes to incentivize this approach, these defense product companies could play an important role in delivering future warfighting capabilities. This paper explores the factors that contributed to the emergence of defense product companies and some of the challenges and opportunities to scale this model.

## Aligning Defense Products to National Security Space Needs

In 2025, Palmer Luckey told *60 Minutes*, “the idea behind Anduril was to build not a defense contractor, but a defense product company” (60 Minutes, 2025). While describing one company's approach, this distinction also signals a broader shift in the way that companies are applying commercial approaches to military innovation. Rather than building solutions with government funding to meet government requirements, this business model relies on self-funded research and development to build systems and technologies that anticipate future warfighting needs. It represents both a departure from the traditional contractor approach and an evolution of more recent defense tech approaches. Whether the United States Space Force can incorporate the speed and innovation that defense product companies offer will significantly impact the future of military space capabilities.

Companies that secure U.S. military contracts to provide goods or services are part of the defense industrial base, but that commonality can mask substantial differences in funding, technology development, and market strategy. Three primary models have emerged.

- **Traditional defense contractors** receive government funding to develop systems according to predetermined military requirements. These companies design and build weapon systems through cost reimbursement or fixed-price contracts, with the government bearing most development costs and risk.
- **Dual-use companies** use private capital to develop technologies primarily for commercial customers, then adapt those technologies for military applications. Because



these companies build products to satisfy a large, diverse customer base, military customers have limited ability to direct or specify technology development.

- **Defense product companies** are a distinct model within the broader defense tech ecosystem. Like dual-use companies, they use private capital to self-fund development. Like traditional contractors, they design exclusively for military customers. Most notably, these companies develop products before government requirements exist, assuming both financial and technical risk while anticipating future military needs.

National security space is uniquely positioned to benefit from the defense product company model. Military space operations are less doctrinally mature than military operations in other domains. While the Space Force has begun defining operational concepts for both offensive and defensive counterspace operations, the service has not yet defined its associated capability and technology requirements. This creates opportunities for companies to anticipate future requirements, enabling product-led innovation. Space also creates overlapping demand for technologies that support commercial and military uses, with dual-use applications ranging from launch and propulsion to guidance and communication. However, counterspace capabilities require technologies that are unique to military applications. This creates a market gap between military and commercial demands, in which military needs are too specialized for dual-use companies but too uncertain for traditional defense contractors. Defense product companies can help bridge this gap by funding and developing innovative solutions to both current and future military space operations.

Defense product companies could serve a critical role in delivering future warfighting capabilities, particularly as the United States continues to explore ways to accelerate the speed of military capability delivery to keep pace with peer competitors. This paper explores the development of defense product companies and their role in military space acquisitions through four parts. First, it examines the evolution of defense product companies, including historical precedents from Northrop's F-20 and General Atomics' Predator. Second, it analyzes current defense product companies, with brief case studies of several leading firms. Third, it evaluates how recent changes to acquisition policy and Space Force planning efforts, including international partner engagement, might shape the future of defense product companies. Fourth, it considers risks to this emerging model, including oversight and integration challenges. The paper concludes with implications for the U.S. Space Force.

## **The Evolution: From Defense Tech to Defense Products**

The emergence of defense product companies is the result of changes in acquisition policy, technology, funding, and warfighting. While traditional defense contractors have long self-funded research and development efforts to create or improve products for military customers, the government may end up reimbursing that spending in future contracts (Lofgren, 2022). Defense product companies represent a unique approach, yet past acquisition efforts were instrumental in pathfinding elements of this new approach. These historical examples help explain how the defense industry has evolved to allow the defense products approach, providing insights into the opportunities and risks that exist.

## **Great Expectations Unfulfilled: The F-20 Tigershark**

Northrop's F-20 Tigershark fighter aircraft was a technically advanced, self-funded development effort designed to address the export market for lightweight fighters at a time when Cold War tensions continued to drive military spending in the United States and around the world. Building on the success of its F-5 program, which delivered a lucrative platform with sales of more than 2500 units to dozens of countries, Northrop invested \$1.2 billion in the F-20 as a product for both the U.S. military as well as U.S. allies and partners (Martin and Schmidt, 1987).



It was also meant to avoid restrictive arms control policies, meaning that its primary competition would be foreign aircraft. However, the Reagan administration relaxed exports controls, putting the F-20 in direct competition with another advanced U.S. aircraft, General Dynamics F-16, for which the U.S. Air Force had already committed funding. Northrop's decision to self-fund the F-20 put it at a disadvantage in securing foreign customers, since foreign military sales of the F-16 lowered its unit cost for the Air Force (Dorr, 2014). Failing to convince either the U.S. government or foreign militaries to purchase the F-20, Northrop was forced to terminate the program and write-off the entire development cost without securing any customers.

The F-20 case illustrates the difficulty of anticipating how changes in government policy can affect demand, both domestically and internationally. It also shows the difficulty of shifting bureaucratic momentum once it has aligned behind a particular solution. Northrop's efforts to self-fund a solution may have reinforced exactly the kind of conservative, risk-adverse behavior that Pentagon has sought to disrupt: "After being burned enough times by the military's uncertainty, major defense companies are reluctant to invest much of their own money to develop new technologies and prefer to develop only what the government says it wants and is willing to pay them to build" (Brose, 2020, 234). This history raises the question of what conditions enable contemporary defense product companies to succeed where Northrop failed.

### **The MQ-1 Predator as the Archetype for Future Defense Products**

The development, fielding, and operational use of General Atomics Predator unmanned aerial vehicle (UAV) for both intelligence, surveillance, and reconnaissance (ISR) and precision strike missions serves as a critical bridge from the traditional defense contractor approach to the rise of defense product companies. The platform's success—both in its assigned missions and in encouraging broader development of unmanned aerial platforms—required the timely convergence of funding, warfare, and policy changes that directly relate to contemporary defense product companies.

**Funding.** Beginning in the 1970s, the military services funded several different UAV programs. Many of these early efforts had only limited success, despite Pentagon spending of more \$2 billion for system development and procurement between 1979 and 1997 (United States General Accounting Office, 1997). A mix of company self-funding and government support—through both the Defense Advanced Research Project Agency (DARPA) and the Central Intelligence Agency—combined to enable development of the Predator UAV's predecessor system, the GNAT 750; however, the company that produced the GNAT declared bankruptcy and sold the technology to General Atomics when it was unable to convert early progress into a major acquisition effort (Strickland, 2013). General Atomics continued self-funding system development because it, "believed in the revolutionary role of these unmanned technologies" (Schneider & Macdonald, 2025, 49). As a result, Predator UAV technology was ready when trends in warfare and policy shifts demanded new solutions.

**Warfare.** The Pentagon's increasing demand for UAVs through the 1990s and 2000s reflected significant changes to warfighting needs and new operational environments. As Schneider and Macdonald (2025) note, "Over the course of the 1990s, the United States opted into a series of lower-intensity conflicts in which policy makers (and defense practitioners) sought to achieve victory with minimal human cost... these low stakes conflicts provided the initial impetus for research, investment, and early experimentation with what would become the MQ-1 Predator unmanned aerial vehicle" (Schneider & Macdonald, 2025). Although beginning as an ISR platform, the Predator was armed with an AGM-114 Hellfire missile for operations in Iraq and Afghanistan after 9/11.

**Policy.** In 1994, the Pentagon used a novel acquisition approach to assess the Predator's utility before committing to large scale procurement. With this approach, the



Advanced Concept Technology Demonstration (ACTD), the Department was able to deploy prototypes of the Predator in Bosnia, providing a realistic assessment of the system's capabilities in an operational context (General Accounting Office, 1999). The ACTD was a specialized tool intended to rapidly field capabilities by demonstrating mature technologies with heavy user involvement and approval throughout the process (Thirtle et al., 1997). Users were able to provide feedback that directly impacted development efforts. The Predator UAV completed development during the 30-month ACTD, with the Air Force assuming operational control of the Predator demonstration assets in 1996 (General Accounting Office, 1997). This history shows the importance of aligning acquisition policy with both user needs and available solutions. For defense product companies, it also underscores the importance of demonstrations and prototyping to build confidence and support among users.

## **Characteristics of Defense Product Companies**

While there is not a single business model or description that encompasses all business models, the companies that are pathfinding this approach share a combination of characteristics that differentiate them from traditional defense contractors, dual-use companies, and other defense tech companies. The primary differentiating characteristics relate to anticipatory research and development, financing and risk allocation, and product design. Defense product companies combine these characteristics in different ways. Moreover, their business models prioritize flexibility to pursue opportunities as they arise. New defense product companies are likely to find different ways to combine these characteristics to suit their product and market.

The primary characteristic of defense product companies is anticipatory research and development. By completing development work ahead of demand, the company can provide a complete or nearly complete product once customer demand materializes. The company can also work with potential customers throughout development to better understand and respond to customers' needs. This can have the added benefit of building customer trust and anticipation, priming the demand that's ultimately needed to close the business case.

The second characteristic is financing and risk allocation. Early-stage defense product companies depend on private capital, including venture capital, to fund initial product development efforts. This private capital enables risk taking and a longer interval before profitability is achieved. Private capital interest in defense tech and defense product companies has grown significantly over the last decade. Prominent venture capitalists have identified this interest as a counter to China's civil-military fusion strategy, specifically by enabling private capital to fund innovative solutions to military problems (Shah & Kirchhoff, 2024). However, since 2021, venture capital funding has declined partly due to government fiscal policy and rising interest rates (Chen et al., 2024). For many companies, access to capital is no longer the key constraint; they must now prove out their business model by achieving profitability (AIAA, 2025).

The third characteristic is product design. A key differentiator between defense product companies and the broader category of defense tech companies is a focus on developing software-enabled hardware, rather than just software. This difference carries several important implications. Most notably, hardware is more capital intensive than software. This creates a higher obstacle to profitability but also a stronger competitive advantage for companies that are successful. Focusing on hardware also expands opportunities to meet military-specific needs. The following case studies of three companies—Anduril, True Anomaly, and Epirus—briefly describe how each has implemented these characteristics into their business models.





## **Anduril**

Since 2019, Anduril has grown significantly, expanding its product offerings into a broad portfolio of capabilities. These products are enabled and integrable within a common software platform, Lattice, that facilitates autonomy, command and control, and data fusion. Venture capitalist and Anduril co-founder Trae Stephens has said that the company's focus on software-enabled hardware was a deliberate strategy to avoid the challenges of selling software to the U.S. military (Lofgren, 2022b). Anduril has raised billions through multiple venture capital-backed funding rounds, leading to a 2025 valuation of just over \$30 billion (Stone & Biswas, 2026). The company's first product for government buyers was the Sentry Tower, an object detection, identification, and tracking system deployed in the United States to support border patrol missions. The company has since created several other discrete products around this core technology, including the Seabed Sentry for persistent and real-time underwater sensing and communications (Anduril, 2026). The company has also expanded across domains, with UAV and space products also enabled by Lattice. Among defense product companies, Anduril exemplifies the broadest business model, building multi-domain capabilities to support a wide range of military users and missions.

## **True Anomaly**

True Anomaly's has developed a business model tightly focused on the Space Force and its military-specific needs. The company's senior vice president for space defense, Stephen Kitay, described True Anomaly as a defense products company able to rapidly and affordably deliver capabilities designed operational concepts for space superiority (Center for Space Policy and Strategy, 2025). Like Anduril, its hardware systems are enabled by a common software platform, Mosaic, supporting both traditional space operations capabilities as well as those required for space control operations. The company's primary hardware product is Jackal, a multi-mission, multi-orbit space vehicle designed for maneuverability and configurability as well as high-volume production. Jackal's capabilities focus on those most applicable for future national security space operations, which co-founder Even Rogers described as, "the small number of things that make space-to-space engagements successful outcomes: sensing, autonomy, maneuverability, and payload-carrying capacity" (Clark, 2026). True Anomaly is a private company and has raised several hundred million dollars through venture capital-backed funding rounds. Its focus on the Space Force as its primary customer presents a challenge, as it will need to tightly couple its development with the service's needs; however, the potential growth of counterspace operations also presents significant opportunities with less entrenched competition than other domains.

## **Epirus**

Epirus designs and manufactures directed energy systems for counter-UAV missions. The company privately funded research, development, and validation of its core system before winning U.S. Army contracts in 2023 and 2025 (Epirus, 2023; Maher, 2025). Like Anduril and True Anomaly, Epirus has raised funding through venture capital, yet it differs in its software approach. While the system allows for improved functionality through software updates, its hardware is designed to integrate within larger software platforms. Notably, Epirus has also identified potential commercial applications for some of its core technology, with potential applications for its radiofrequency technology to support commercial communications and for its power management technology to support commercial energy applications.

## **Converting the Demand Signal into Action**

Since 2023, the Ronald Reagan Institute has published an annual "report card" for the National Security Innovation Base. The reports' grading system assesses both government and industry reforms, efforts, and initiatives to strengthen the United States' innovation and



technological edge. The reports generally show relatively small changes from year-to-year (e.g., improvements from “B” to “B+”). But in 2026, the Ronald Reagan Institute recorded its largest ever grade improvement, moving from a “D” to a “B-”, in its assessment of “customer clarity,” which the organization defines as the “demand signal for customer (government) innovation priorities, including funding and acquisition pathways to match the aspiration” (Ronald Reagan Institute, 2026).

Greater clarity has not substantially affected prime contractors’ dominance of Department of War contracts. The Reagan Institute’s 2026 report notes that while the Pentagon increasingly obligated funding to defense tech companies—both in relative and absolute terms—between fiscal years 2023 and 2025, those companies still receive less than 1% of total obligations. A 2025 Defense Innovation Board identified technical debt as one of the factors contributing to this mismatch between the Pentagon’s stated desire to incorporate new defense technology companies and its limited success to date (Defense Innovation Board, 2025). Specifically, many of these companies focus on advanced, software-intensive solutions that cannot be readily integrated within the Department’s outdated, legacy systems. A similar problem confronts military hardware acquisitions. Analyses of recent U.S. military action against Iran, Operation Epic Fury, have noted that nearly all the weapons the United States has used to date are decades old (Hoff, 2026). The Pentagon is requesting supplemental funding, largely to replenish munitions for these legacy systems.

The slow pace of change in Pentagon spending on defense technology and defense product companies is a product of many factors. Technical debt, overreliance on legacy platforms and systems, the software-intensive nature of some nontraditional vendors, and the relative immaturity of some defense product companies. Part of these companies’ business case is developing products to meet future demand. Thus, there is a lag between product development efforts and when government funding obligations should reasonably be expected, which partly explains the relatively high valuations of some of these companies compared to their government-related revenues. Military leaders, servicemembers, and Congress face strong disincentives to trying novel solutions or new approaches to deliver military capabilities, particularly for emergent threats that struggle to command the same attention as today’s challenges (Brose, 2020). To be a better buyer, the government must avoid defaulting to large development programs when current commercial technologies do not meet immediate needs (AIAA, 2025). The more appropriate comparison is where commercial technology will be when the development program is complete, typically five or more years in the future. The government can then help inform commercial innovation by clarifying its needs and anticipated demand. Solutions will have to be supported by policy and cultural change that promote informed risk-taking with a realistic assessment of both future needs and the rate of innovation among defense product companies

## **Acquisition Transformation and Space Force’s Future Force Structure**

### **Acquisition Transformation**

In 2025, the Pentagon announced a series of acquisition reform efforts that carry long-term implications for companies across the defense industrial base. Between August and November, Secretary Hegseth signed memos that announced major changes to the processes that govern military requirements, acquisitions, and related activities. These changes were intended to reflect the Department’s shift to a more assertive warfighting posture, while also addressing longstanding, systemic deficiencies that had contributed to schedule delays, cost overruns, and erosion of the United States’ military advantage. A critical step to achieving these goals is expanding the number and diversity of companies in the defense industrial base: “the Department will restore innovation, accelerate production, gain negotiating leverage, and more



effectively manage cost and schedule growth by promoting competition with lowered barriers to entry and diversification of prime and subcontractor sourcing without sacrificing quality” (Office of the Under Secretary of War for Acquisition and Sustainment, 2025). If successfully implemented, several of these reforms would benefit defense product companies seeking to capture military contracts and validate their business models. Three reforms are particularly relevant for defense product companies: requirements validation reform; creation of portfolio acquisition executives (PAE); and policies to encourage self-funded research and development and increased production rates.

**Requirements Validation.** The first reform memo, signed in August 2025, disestablished the existing Joint Capabilities Integration and Development System (JCIDS), shifting requirements validation authority from the Joint Staff to the military services (Hegseth & Feinberg, 2025). The reform is intended to streamline reviews and shorten the timeline between when a capability need is identified and validated. The slow pace of requirements validation has been a central criticism of the acquisition system because it delays development work to close a capability gap. But even for a defense product company that has little to no additional development work planned, a slow requirements validation process delays contract awards, and with it the company’s revenue to fund future products.

**Portfolio Acquisition Executive.** The second reform memo, signed in November 2025, initiates a broad overhaul of the acquisition system, establishing rapid capability delivery as the Department’s organizing principle for acquisitions (Hegseth, 2025). A key mechanism for accomplishing this goal is redesignating program executive officers (PEO) as portfolio acquisition executives (PAE) with expanded authority.

A perennial criticism of the PEO construct is that it aligns incentives around programs rather than capabilities. Because programs are difficult both to start and to change once underway, acquisition leaders are disincentivized to consider novel solutions or abandon a chronically underperforming program. In contrast, the PAE construct centers incentives around delivering a portfolio of related capabilities through multiple programs, rather than optimizing for individual programs. This both enables and incentivizes PAEs to execute cost, schedule and performance trades across acquisition programs. PAEs also have greater control over functional support—contracting, finance, systems engineering, etc.—to make decisions across the portfolio.

Portfolio management enables PAEs to strengthen competition by disaggregating capability requirements and combining solutions from different companies. This creates opportunities for defense product companies to provide specific solutions within capability portfolios and to compete for contracts even after losing initial development or production awards. Incumbents remain incentivized to maintain high performance because contract awards do not insulate them from competition for delivering the system or capability.

**Self-funded Research and Development and Increased Production Rates.** The Department’s reforms include directives aimed at encouraging companies to increase their research and development spending and their production rates. While specific policies will be largely implemented through Department- and service-level guidance, the November 2025 memo identifies both strategic direction and specific policy changes. Defense product companies, having already invested private capital in military solutions, hold a substantial advantage over companies that wait for government funding before starting development activities.

At the strategic level, the reforms establish an acquisition priority for commercial products and offerings, defined broadly as “industry solutions funded by private investment that meet military needs” (Exec Order No. 14265, 2025). This definition encompasses both dual-use



commercial products and defense product companies' military specific solutions. The Department's Acquisition Transformation Strategy further clarifies that stabilizing demand signals and correctly aligning incentives will encourage industry investment (Office of the Under Secretary of War for Acquisition and Sustainment, 2025). In other words, by rewarding companies that self-fund innovation and capacity, the Department aims to drive a virtuous cycle of further private investment.

### Space Force Planning for Emerging Missions

Since the Space Force was created in 2019, the service has been evolving its strategic planning processes to inform long-term capability goals. Most of the Space Force's current acquisition programs are either inherited from the Air Force or are previously planned replacements for existing systems. However, the combination of increasing budget authority, emerging threats, and expanding missions is now empowering the Space Force to adopt proactive planning approaches, considering both the future threats Guardians will confront and the capabilities they will need to counter those threats. With the Department's focus on speed to capability delivery and scale, the Space Force's efforts to design and communicate its future force structure become an essential acquisition function. These planning efforts provide critical signals of the service's needs and priorities, assisting industry in aligning its research and development activities.

U.S. policy has long rejected the common perception that space is an operational "sanctuary" (Dickey, 2020), the Space Force has, only recently, begun openly discussing the role of offensive space capability in supporting U.S. military objectives. In March 2025, the Space Force released its *Space Warfighting* framework, "[establishing] basic principles for the use of military power" to achieve space superiority, which includes both offensive and defense actions under the USSF "core function" of space control (United States Space Force, 2025). Figure 1 lays out the relationship between these terms and the categories of offensive and defensive action, as outlined by the framework.




**Figure 1. Space Warfighting Terminology**

Observers have noted that this terminology place space operations on a similar lexical plane as other domains, for which strategists and practitioners have long used concepts of superiority and control to describe their objectives (Hadley, 2025). Detailing the actions that compose counterspace operations also reveals similarities and differences between military and commercial operations that separate defense product and dual-use companies. Table 1 shows the full list of counterspace actions, by category, identified in the strategy.

**Table 1. Military-Specific Actions in Counterspace Operations**

| <b>Offensive</b>                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>Defensive</b>                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Orbital strike</b> <ul style="list-style-type: none"> <li>• Pursuit</li> <li>• Standoff</li> </ul> <b>Space link interdiction</b> <ul style="list-style-type: none"> <li>• Electromagnetic attack</li> <li>• Cybernetwork attack</li> </ul> <b>Terrestrial strike</b> <ul style="list-style-type: none"> <li>• Ground-based fires</li> <li>• Maritime-based fires</li> <li>• Air-based fires</li> <li>• Space-based fires</li> </ul> | <b>Active space defense</b> <ul style="list-style-type: none"> <li>• Escort</li> <li>• <b>Counterattack</b></li> <li>• Suppression of adversary counterspace targeting</li> </ul> <b>Passive space defense</b> <ul style="list-style-type: none"> <li>• Threat warning</li> <li>• <b>Military deception</b></li> <li>• Hardening</li> <li>• Dispersal</li> <li>• Disaggregation</li> <li>• Mobility</li> <li>• Redundancy</li> </ul> |

*Note:* Underlying information in the table is reproduced from the U.S. Space Force's *Space Warfighting* framework. **Bold** indicates that author's judgement of military-specific actions that lack a clear commercial application.

A more thorough analysis of these actions is beyond the scope of this paper; however, even a cursory review of these categories reveals important considerations for both defense product and dual-use companies. First, and unsurprisingly, explicitly offensive actions within counterspace operations lack a commercial application. Second, most defensive actions within counterspace operations have at least a potential commercial application. Because space is a distant and harsh environment that can be congested with both operational systems and debris, commercial space systems employ features like disaggregation, mobility, and redundancy to protect their investments and maintain services. Even less clear-cut examples, like escort and suppression of counterspace targeting are feasible given concerns that space systems may be targeted by malicious actors for criminal, military, or other purposes. However, the specific application of a defensive action may tip into a military-specific use case. For example, commercial companies may harden their satellites against radiation from space weather events, while the military will harden some satellites against radiation from a nuclear detonation.

Further refinement of these categories and innovation in related technologies can help further differentiate the boundaries that separate solutions from defense product and dual-use companies. As the Space Force continues to mature its future force structure and long-term operational goals, it can signal additional demands that are distinct from those of commercial space companies. In turn, those signals can help inform private investment in military-specific capabilities.





## **Addressing Challenges to the Defense Product Model**

### **Oversight**

As defense product companies become more prominent, the Space Force will need to ensure that effective oversight is able to evolve in ways that accommodate this distinct business model. For traditional acquisition programs, substantial oversight—whether conducted by the service, the Department, or Congress—occurs while a system is under development through program reviews, technical assessments, and cost and schedule metrics. During this period, oversight bodies monitor as the system design is refined through experimentation and prototyping, helping ensure the final system will meet performance requirements as well as cost and schedule commitments before proceeding to production. This oversight is intended to ensure agreement among stakeholders over the program’s goals and its progress against those goals while there is still flexibility to course correct as needed.

Defense product companies represent a challenge to the traditional oversight model because development work occurs before government involvement. Two specific oversight challenges merit attention. First, this timing means that the military customer, as well as oversight bodies, may have limited insight into design decisions and technical approaches that shape product performance, requiring alternative mechanisms to validate that the product will perform as claimed in operational conditions. Defense product companies will need to build and maintain trust as well as provide data or other means of allowing independent verification of the company’s claims. Second, intentionally or unintentionally, requirements and contracts may be written in such a way to unduly favor a particular company’s offering over competing solutions. The Space Force will need to ensure that solicitations focus on broad capability requirements rather than product-specific features. The service may also need to explore new oversight mechanisms, balancing acquisition speed with accountability. Equally important, oversight bodies will need to adjust their expectations and practices to avoid smothering innovation with traditional methods.

### **Integrated Capabilities and Competition**

The Space Force recognizes that future warfighting depends on fielding integrated capabilities, with space systems as integral to the kill chains that enable the joint force’s effectiveness (Space Training and Readiness Command, 2025). National security space systems must efficiently generate and relay data across multiple networks and systems to execute military effects on complex, dynamic battlefields. Systems designed around proprietary or vendor-specific architectures—even if highly integrated within that network—present both operational and acquisition challenges by limiting interoperability.

Defense product companies may face a conflict reconciling integration with their business models. Because many build capabilities and value through software, they may be less incentivized to support integration with competing vendors’ systems. For some of the largest defense product companies, the business model depends on securing long-term service or sustainment contracts by providing foundational software capabilities that become embedded in military operations (Panter, 2025). Over the near term, this strategy may work against the Space Force’s efforts to improve integration and interoperability.

The Army’s 2026 contract with Anduril illustrates potential challenges with single-vendor platform approaches. The Army awarded Anduril a 10-year, fixed price contract with a ceiling up to \$20 billion to consolidate current and future commercial solutions into a single, enterprise licensing agreement. While this could simplify the Army’s access to commercial technology and ensure consistent pricing and terms, analysts have noted that this approach could also create dependency on Anduril’s proprietary software at the integration layer (Panter, 2025). According to some experts, the Army may also become complacent in reaching for familiar solutions from



established vendors rather than seeking out more disruptive innovations (Pomerleau & Welch, 2026). Additionally, this approach may serve as a false demand signal for industry if not followed by consistent orders and stable budgets. Both challenges are likely to appear in the government's agreements with defense product companies.

Within space acquisitions, the Space Development Agency's experience delivering its Proliferated Warfighting Space Architecture (PWSA) illustrates multi-vendor integration challenges. SDA prioritized competition through regular contract awards split among multiple vendors to field several hundred satellites (Berglund, 2024). However, according to SDA officials, integrating satellites from multiple vendors has proved challenging, contributing to delays and requiring a dedicated integration support contract that was not part of the PWSA's initial acquisition strategy (Erwin, 2025). While the PWSA involves both traditional contractors and dual-use companies, similar challenges would need to be addressed to combine solutions from multiple defense product companies to deliver an integrated space capability.

The Space Force will need to establish clear integration and interoperability requirements in contracts with defense product companies, while maintaining competitive pressure to prevent vendor lock-in when possible. This will require a variety of mechanisms, including technical standards, open architecture requirements, and contract structures that balance platform integration and scale with competition.

### **Market limits and Exportability**

Defense product companies confront a structural challenge inherent to their business model: reliance on a single, volatile customer. Repeated studies have documented how the volatility of U.S. military demand—whether driven by budget cycles, political shifts, or system recapitalization schedules—is an obstacle for companies trying to maintain stable production lines, supply chains, and costs (Section 809 Panel, 2019). Even if a defense product company secures contracts from multiple government agencies or military services, it still faces uncertainty inherent to U.S. government procurement.

International sales to allies and partners are one mechanism for mitigating government demand volatility. In March 2026, Anduril completed delivery of nearly 300 of its Altius drone, a versatile platform that can support multiple missions, and supporting capabilities to Taiwan as part of a foreign military sales package the United States approved in 2024 (Lo et al., 2026). According to Anduril, the company was able to rapidly deliver on the contract by self-funding production and are taking additional steps to establish a long-term partnership with the Taiwanese military and industrial base (Anduril, 2025).

Similar foreign opportunities may be possible for defense product companies that focus on space solutions. While those companies may face greater export scrutiny due to the sensitivity of the counterspace missions they are likely to support, sales to close U.S. allies with sophisticated space capabilities may be acceptable. Such sales would also align with ongoing Space Force efforts to better integrate the space capabilities of allies and partners. For example, the Space Force's 2025 International Partnerships Strategy initiated a series of activities to better include allies and partners in future force planning (United States Space Force, 2025b).

### **Conclusion: Implications for USSF Acquisition**

The United States faces a challenging geopolitical and strategic environment, with mounting pressure from peer adversaries. The Space Force will need contributions from traditional defense contractors, dual-use companies, and defense product companies to confront challenges in the space domain. A clearer definition of its strategic objectives and capability requirements, particularly offensive and defensive actions within counterspace



operations, the service is sending a clearer signal to all DIB companies to inform research and development efforts. National security space will continue to benefit from commercial innovations for its many dual-use applications, but it will increasingly also demand military-specific capabilities as the Pentagon pursues new missions in space. The coming years represent an important test for both the Space Force and the broader defense acquisition community to demonstrate progress seizing commercial innovation to rapidly field warfighting capabilities.

*Approved for public release. OTR 2026-00568.*

## References

- 60 Minutes. (2025, May 18). "Palmer Luckey on making autonomous weapons for the U.S. and its allies," [video]. YouTube. <https://youtu.be/bWEXnph1EI?si=3vufLdLmeWeE4tmz>
- AIAA. (July 24, 2025) *Jamie Morin Speaks on Savvy Space Policy, 2025 AIAA David W. Thompson Lecture* [Video]. YouTube. <https://www.youtube.com/watch?v=HEq6rHjmw2Y>
- Anduril. (2026). *Seabed Sentry*. <https://www.anduril.com/seabed-sentry>
- Anduril Industries. (2025, August 5). *Anduril Expands Commitment to Taiwan with New Capabilities, Partnerships and Local Investment*. Anduril. <https://www.anduril.com/news/anduril-expands-commitment-to-taiwan-with-new-capabilities-partnerships-and-local-investment>
- Berglund, A. (2024, July 25). *The Space Development Agency and the Future of Defense Space Acquisitions*. Center for Space Policy and Strategy. [https://csps.aerospace.org/sites/default/files/2024-07/Berglund\\_SDA\\_20240716.pdf](https://csps.aerospace.org/sites/default/files/2024-07/Berglund_SDA_20240716.pdf)
- Berkenstock, D., & Philips, H. (2025). *The Defense Tech Playbook: A Strategic Planning Resources for Venture-Backed Startups Building for the US Defense Community*. The Hoover Institution. [https://www.hoover.org/sites/default/files/research/docs/Berkenstock\\_DefenseTechPlaybook\\_web-250213.pdf](https://www.hoover.org/sites/default/files/research/docs/Berkenstock_DefenseTechPlaybook_web-250213.pdf)
- Brose, C. (2020). *The Kill Chain: Defending America in the Future of High-Tech Warfare*. Hachette Books.
- Center for Space Policy and Strategy. (2025, July 23). *Rapid Acquisition for Defense*. The Space Policy Show, Episode 158 [Video]. YouTube. <https://www.youtube.com/watch?v=31BUa8QWcmg>
- Chen, W.W., Han, M.R., Georgin, S.E., and Wilson, R.S. (2024, October 24). *Leverage and Preserve: The Need for DOD to Strengthen Support for U.S. Commercial Space*. Center for Space Policy and Strategy. <https://csps.aerospace.org/papers/leverage-and-preserve-need-dod-strengthen-support-us-commercial-space>
- Clark, S. (2026, March 19). *Dogfighting in space won't look like the movies, but this company wants in on it*. Ars Technica. <https://arstechnica.com/space/2026/03/dogfighting-in-space-this-company-builds-satellites-for-high-tempo-engagement/>
- Defense Innovation Board. (2025). *Scaling Nontraditional Defense Innovation*. Department of Defense. [https://stib.cto.mil/wp-content/uploads/2026/01/2025-2\\_DIB-ScalingNontraditionalDefenseInnovation\\_250113PUBLISHED\\_9ee4ae.pdf](https://stib.cto.mil/wp-content/uploads/2026/01/2025-2_DIB-ScalingNontraditionalDefenseInnovation_250113PUBLISHED_9ee4ae.pdf)
- Dickey, R. (2020, September 1). *The Rise and Fall of Space Sanctuary in U.S. Policy*, Center for Space Policy and Strategy. <https://csps.aerospace.org/papers/rise-and-fall-space-sanctuary-us-policy>
- Dorr, R.F. (2014, August 29), *What Might Have Been: F-20 Tigershark*. DefenseMediaNetwork. <https://www.defensemedianetwork.com/stories/what-might-have-been-f-20-tigershark/>



- Epirus. (2023, January 23). *U.S. Army Awards Epirus \$66.1M Contract for Leonides Directed Energy System*. <https://www.epirusinc.com/press-releases/u-s-army-awards-epirus-66-1m-contract-for-leonidas-tm-directed-energy-system>
- Erwin, S. (2025, March 12). Space Development Agency Adjusts Satellite Procurement Strategy. *Space News*. <https://spacenews.com/space-development-agency-adjusts-satellite-procurement-strategy/>
- Hadley, G. (2025, April 21). New Doc Spells Out How USSF Will Use Space Control to Gain Space Superiority. *Air & Space Forces Magazine*. <https://www.airandspaceforces.com/new-doc-spells-out-how-ussf-will-use-space-control-to-gain-space-superiority/>
- Hegseth, P., & Feinberg, S. (2025). *Reforming the Joint Requirements Process to Accelerate Fielding of Warfighting Capabilities*. U.S. Department of Defense. <https://media.defense.gov/2025/Nov/10/2003819442/-1/-1/1/REFORMING-THE-JOINT-REQUIREMENTS-PROCESS-TO-ACCELERATE-FIELDING-OF-WARFIGHTING-CAPABILITIES.PDF>
- Hegseth, P. (2025). *Transforming the Defense Acquisition System into the Warfighting Acquisition System to Accelerate Fielding of Urgently Needed Capabilities to Our Warriors*. U.S. Department of Defense. <https://media.defense.gov/2025/Nov/10/2003819439/-1/-1/1/TRANSFORMING-THE-DEFENSE-ACQUISITION-SYSTEM-INTO-THE-WARFIGHTING-ACQUISITION-SYSTEM-TO-ACCELERATE-FIELDING-OF-URGENTLY-NEEDED-CAPABILITIES-TO-OUR-WARRIORS.PDF>
- Hoff, R. (2026, April 6). *America Can't Win Tomorrow's Wars with Yesterday's Defense Budget*. Real Clear Defense. [https://www.realcleardefense.com/articles/2026/04/06/america\\_cant\\_win\\_tomorrows\\_wars\\_with\\_yesterdays\\_defense\\_budget\\_1174861.html](https://www.realcleardefense.com/articles/2026/04/06/america_cant_win_tomorrows_wars_with_yesterdays_defense_budget_1174861.html)
- Lo, T. P., Wu, C. Y., & Chin, J. (2026, March 20). US completes deliveries of suicide drones. *Taipei Times*. <https://www.taipetimes.com/News/taiwan/archives/2026/03/20/2003854149>
- Lofgren, E. (2022, May 18). *Defense Primes Have an Advantage in OTA Dues to a Quirk in IRAD*, Acquisition Talk. <https://acquisitiontalk.com/2022/05/defense-primes-have-an-advantage-in-otas-due-to-a-quirk-in-irad/>
- Lofgren, E. (2022b, July 14). *The Trojan Horse: Wrapping Software in Metal to Make the Sale to DOD*, Acquisition Talk. <https://acquisitiontalk.com/2022/07/the-trojan-horse-wrapping-software-in-metal-to-make-the-sale-to-dod/>
- Maher, T. (2025, July 17). *Army awarding Epirus 43.5M for second-generation IFPC high-power microwave systems*. Inside Defense. <https://insidedefense.com/daily-news/army-awarding-epirus-435m-second-generation-ifpc-high-power-microwave-systems>
- Martin, T., & Schmidt, R., (1987, June) *A Case Study of the F-20 Tigershark*. RAND Corporation. P-7495-RGS.
- Office of the Under Secretary of War for Acquisition and Sustainment. (2025). *Acquisition Transformation Strategy*. U.S. Department of War. <https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/ACQUISITION-TRANSFORMATION-STRATEGY.PDF>
- Panter, J. (2025, August 26). *A Skeptic's View of the Hype Machine and Business Model of Neo-Defense Tech*. Cogs of War. <https://warontherocks.com/2025/08/a-skeptics-view-of-the-hype-machine-and-business-model-of-neo-defense-tech/>



- Pomerleau, M., & Welch, C. (2026, April 2). *Army's enterprise contracting vehicles likely to speed procurement but not without risk*. Breaking Defense. <https://breakingdefense.com/2026/04/armys-enterprise-contracting-vehicles-likely-to-speed-procurement-but-not-without-risk/>
- Ronald Reagan Institute. (2026). *2026 National Security Innovation Base Report Card*. [https://www.reaganfoundation.org/cms/assets/1773938920-2026\\_nsib\\_report-card.pdf](https://www.reaganfoundation.org/cms/assets/1773938920-2026_nsib_report-card.pdf)
- Schneider, J., & Macdonald, J. (2025). *The Hand Behind Unmanned: Origins of the U.S. Autonomous Military Arsenal*. Oxford University Press.
- Section 809 Panel. (2019). *Report of the Advisory Panel on Streamlining and Codifying Acquisition Regulations: Volume 3*. [https://discover.dtic.mil/wp-content/uploads/809-Panel-2019/Volume3/Sec809Panel\\_Vol3-Report\\_Jan2019\\_part-1\\_0509.pdf](https://discover.dtic.mil/wp-content/uploads/809-Panel-2019/Volume3/Sec809Panel_Vol3-Report_Jan2019_part-1_0509.pdf)
- Shah, R. M., & Kirchhoff, C. (2024). *Unit X: How the Pentagon and Silicon Valley are Transforming the Future of War*. Scribner.
- Space Training and Readiness Command. (2025). *Space Force Doctrine Document – 1: The Space Force*. United States Space Force. [https://csps.aerospace.org/sites/default/files/2025-04/Space%20Force%20Doctrine%20Document%201%20%28SFDD-1%29\\_4Apr25\\_0.pdf](https://csps.aerospace.org/sites/default/files/2025-04/Space%20Force%20Doctrine%20Document%201%20%28SFDD-1%29_4Apr25_0.pdf)
- Strickland, Frank. (2013). The Early Evolution of the Predator Drone. *Studies in Intelligence*, 57(1). <https://www.cia.gov/resources/csi/static/Early-Evolution-of-Predator.pdf>
- Stone, M., & Biswas, P. (2026, March 3). *US defense firm Anduril set to double its valuation with \$4 billion funding, source says*. Reuters. <https://www.reuters.com/business/anduril-set-double-its-valuation-with-4-billion-funding-round-source-says-2026-03-03/>
- Thirtle, M.R., Johnson, R.V., Birkler, J.L. (1997). *The Predator ACTD: A Case Study for Transition Planning to the Formal Acquisition Process*. RAND Corporation. [https://www.rand.org/pubs/monograph\\_reports/MR899.html](https://www.rand.org/pubs/monograph_reports/MR899.html)
- United States General Accounting Office. (1999). *Unmanned Aerial Vehicles: DOD's Demonstration Approach has Improved Project Outcomes*. <https://www.gao.gov/assets/nsiad-99-33.pdf>
- United States General Accounting Office. (1997) *Unmanned Aerial Vehicles: DOD's Acquisition Efforts*. <https://www.gao.gov/assets/t-nsiad-97-138.pdf>
- United States Space Force. (2025a). *Space Warfighting: A Framework for Planners*. [https://csps.aerospace.org/sites/default/files/2025-10/USSF%20Space\\_Warfighting\\_A%20Framework%20for%20Planners%20%28final\\_20250410%29.pdf](https://csps.aerospace.org/sites/default/files/2025-10/USSF%20Space_Warfighting_A%20Framework%20for%20Planners%20%28final_20250410%29.pdf)
- United States Space Force. (2025b). *U.S. Space Force International Partnership Strategy: Strength Through Partnership*. <https://csps.aerospace.org/sites/default/files/2025-10/USSF%20International%20Partnership%20Strategy%20%28June%202025%29.pdf>





## PANEL 6. GUARDING THE MACHINE: PRESERVING HUMAN AGENCY AND MARKET INTEGRITY IN THE AI ERA

Wednesday, May 6, 2026, 1315 – 1430 ET (1015 - 1130 PT)

### Panel Summary:

As the Department of War accelerates the adoption of emerging technologies, it faces a complex landscape where autonomous AI algorithms and transition barriers threaten to outpace traditional governance. This panel explores the critical intersection of market integrity and technology readiness, examining how AI-driven collusion in procurement auctions and the rise of "AI-to-AI" proposal ecosystems necessitate a renewed focus on meaningful human oversight. By introducing practical tools like the Technology Maturity Framework (TMaF), these researchers provide a roadmap for maintaining ethical discernment, accountability, and strategic trust while delivering advanced capabilities to the warfighter.

**Chair: CAPT Michael Owen, USN**, Vice Provost for Warfare Studies and Director, Naval Postgraduate School AI Task Force

### Presenters:

**Keeping Humans in Command: Preserving Agency in AI-Driven Federal Acquisition – Adam Bouffard, The MITRE Corporation**

**AI Collusion in Procurement Auctions – Andrew Tai, Defense Resources Management Institute**

**An Actionable Framework & Interactive Web Application for Transitioning Technology to the Warfighter: The Transition Maturity Framework (TMaF) and Tech2PEO/PAE – Alex Cocco, MITRE**



**CAPT Mike Owen**—grew up in central Virginia and graduated with distinction from Virginia Tech in 1997 with a Bachelor of Science in Mechanical Engineering. He also holds a Master of Engineering Management from Old Dominion University and a Master of Arts in Government from Regent University. He is a Joint Qualified Officer and a qualified Joint Targeting Officer.

CAPT Owen began his career as a submarine officer with at-sea assignments in USS MINNEAPOLIS-ST PAUL and USS NORFOLK and as the Submarine Special Projects Officer on the COMUSNAVEUR staff in London, England. Following his re-designation to Naval Intelligence in 2005, his Intelligence assignments have included the Joint Transformation Command-Intelligence as the Head of the Intelligence Quick Reaction Team, SIO on USS

KEARSARGE (LHD-3), Fleet Intelligence Officer at COMSUBLANT, N2 at Naval Special Warfare Group Four (NSWG-4), Director of Intelligence for COMSUBLANT and JFMCC-STRAT, and Director of Operations and Training for Naval Information Warfighting Development Center. Additionally, CAPT Owen deployed to Ramadi, Iraq in 2008 as the Deputy J2 for Task Force Ramadi as part of a team from the JFCOM Joint Enabling Capabilities Command (JECC).

CAPT Owen lives in Monterey, CA with his wife of 26 years, Carrie. They have two adult children, Devyn, 25, and Aiden, 22.



## **Keeping Humans in Command: Preserving Human Agency in AI-Driven Federal Acquisition**

**Stephen W. Roe**—is a Managing Director at MITRE, leading Veterans and Military Health Modernization and supporting VA, VACO, and VHA since 2020. With over 20 years of experience driving cross-agency innovation, he has delivered measurable outcomes across healthcare, acquisition reform, and data-driven policy. His achievements include the Whole Health Recognition Framework (2+ million Veterans), PolicyNet (40% reduction, \$55 million savings), and GenHub (2,000+ studies, 1,500 publications). He has shaped major VA priorities such as the PACT Act, EHR modernization, and community care reform, impacting \$100+ billion in procurement while advancing AI-enabled patient safety and connected care. He holds degrees from Virginia Tech, Johns Hopkins, and UVA Darden. [[sroe@mitre.org](mailto:sroe@mitre.org)]

**Dr. Nathan Turnipseed**—serves as Director of the Acquisition Integration Service within the Department of Veterans Affairs Strategic Acquisition Center, where he provides strategic counsel to senior executives on enterprise acquisition life-cycle management, data governance, and complex acquisition strategy across a multi-billion-dollar portfolio. A veteran with more than 25 years of leadership in public and private sector acquisition environments, he is recognized for advancing large-scale, technology-enabled transformation initiatives and driving organizational performance. His distinguished U.S. Air Force career included leadership in personnel development for enlisted Airmen and commissioned officers. Dr. Turnipseed holds a Doctor of Management in Organizational Leadership, a Master of Management, and a Bachelor of Science in Psychology and Sociology, and maintains a Level III Federal Acquisition Certification for Program and Project Managers (FAC-P/PM). [[Nathan.Turnipseed2@va.gov](mailto:Nathan.Turnipseed2@va.gov)]

**Ryan Novak**—Acquisition Outcome Lead at MITRE, brings 30 years of experience as a USAF Contracting Officer and acquisition innovator. He holds multiple advanced degrees, including two MBAs, and is DAWIA III Contracting certified. A published author and speaker, Ryan created the Challenge-Based Acquisition Handbook and taught innovative acquisition approaches to 800+ professionals across government, industry, and MITRE. He has led strategic programs such as iCAMS, held an unlimited CO warrant, and completed executive programs at Oxford and Yale. As MITRE's Acquisition Innovation Center Lead for AI in Acquisition, he drives AI strategies, tools, and prototypes to transform federal procurement under the ACQAIRES Portfolio. [[rnovak@mitre.org](mailto:rnovak@mitre.org)]

**Adam Bouffard**—is a Group Leader and Principal Decision Analyst at MITRE, with over 20 years of experience in acquisition and contracts supporting multiple DoD and civilian government agencies. Prior to joining MITRE in 2015, he worked as a civilian supporting the Naval Sea Systems Command and the Office of Naval Research. He holds a master's degree in systems engineering from George Washington University and possesses multiple contracts certifications. [[abouffard@mitre.org](mailto:abouffard@mitre.org)]

**Wilson Miles**—is an Associate Research Fellow at ETI. His research and analysis portfolio focuses on emerging technology supply chains, acquisition policy, artificial intelligence, workforce issues, and other modernization technology policy issues. Wilson previously interned at several nonprofit organizations, including CRDF Global, the Hudson Institute, and the Foundation for Defense of Democracies. He earned his master's degree in international affairs: U.S. Foreign Policy and National Security from American University's School of International Service and a bachelor's degree in international relations from Linfield University. [[wmiles@ndia.org](mailto:wmiles@ndia.org)]

**Adam Kitay**—is a Strategic Business Operations Principal at MITRE, with over 20 years of experience in government contracts and acquisition supporting various intelligence, DoD, and civilian agencies. Prior to joining MITRE, Adam worked as a civilian for the Department of the Navy, Naval Air Warfare Center (JSF PMO), the Environmental Protection Agency, and the Missile Defense Agency. He was also a Sr. Manager at Lockheed Martin Space, working in their Special Programs division. Adam holds a Bachelor of Science in Business Administration from Colorado State University. [[akitay@mitre.org](mailto:akitay@mitre.org)]

**Kevin Forbes**—is a Principal Software Engineer at the MITRE Corporation, with over 25 years of software development experience. His current work centers on applying Generative AI to tackle challenges in government acquisition and the software development life cycle. Kevin earned his



bachelor's degree in computer science from Georgetown University and a master's degree in systems engineering from Johns Hopkins University. [[forbes@mitre.org](mailto:forbes@mitre.org)]

**Christopher Barlow**—is an Innovation Strategist with The Belmont Group, with 10+ years of experience and demonstrated expertise in audit risk mitigation, acquisition data analytics, process development and optimization, and innovative acquisition approaches, including challenge-based acquisitions (ChBA) and artificial intelligence applications. His diverse experience spans multiple federal agencies. Chris holds a bachelor's degree in business administration from the University of Pikeville and is a Certified Business Intelligence Professional from Villanova University [[cbarlow@blmtgrp.com](mailto:cbarlow@blmtgrp.com)]

## Abstract

As artificial intelligence (AI) rapidly transforms government operations and private industry, the federal acquisition process faces a new paradox: the increasing use of AI to generate, evaluate, and award contract proposals. This convergence creates an “elephant in the room” scenario, where AI-enabled procurement systems in government may ultimately review proposals drafted by AI tools in industry, potentially diminishing the role of human judgment in critical decision-making processes. While automation offers efficiency, consistency, and scalability, the absence of human oversight risks eroding accountability, fairness, ethical discernment, and institutional knowledge in the acquisition process.

This paper examines the dual-use evolution of AI in federal procurement, specifically how government acquisition offices are adopting AI-driven tools for market research, requirements generation, and proposal evaluation, while industry leverages similar technologies to optimize proposal writing and cost modeling. Drawing on guidance from America's AI Action Plan and other federal AI policy frameworks, this research identifies the core risks of an AI-to-AI procurement ecosystem: diminished transparency, amplified algorithmic bias, and the potential loss of critical human judgment in evaluating qualitative and context-sensitive factors.

The expected outcome of this research is a framework for preserving meaningful human oversight within AI-augmented acquisition environments. This includes recommendations for: (1) defining human-in-the-loop thresholds across acquisition phases; (2) establishing verification and audit mechanisms for AI-generated content; and (3) integrating AI literacy training for acquisition professionals.

By confronting the AI-to-AI dilemma directly, this study seeks to confirm that automation enhances, rather than replaces, the human expertise, ethical reasoning, and strategic judgment that underpin the integrity of federal acquisitions. The goal is a balanced approach where AI supports data-driven decision-making, while humans remain the final arbiters of fairness, accountability, and mission alignment, preserving the essential role of human agency at the core of accelerating warfighting capabilities. This research directly aligns with the theme of this year's Naval Postgraduate School Acquisition Research Symposium, “Accelerating Warfighting Capabilities,” by advancing an ethical and operational framework that confirms AI adoption strengthens, rather than undermines, the speed, trust, and integrity of defense acquisition.

## Introduction

Artificial intelligence (AI) is rapidly reshaping the federal acquisition environment. Government acquisition organizations are beginning to deploy AI-enabled tools to assist with market research, requirements development, cost analysis, and proposal evaluation. At the same time, contractors across the defense industrial base are adopting generative AI and advanced analytics to automate proposal writing, compliance checks, and pricing optimization. These parallel developments are converging within the same procurement ecosystem, fundamentally altering how acquisition information is produced, analyzed, and acted upon.

This convergence introduces a structural governance challenge for federal procurement. When government AI systems evaluate proposals that have been drafted or optimized by AI tools, acquisition decisions may increasingly be shaped by machine-mediated interactions



rather than expert human judgment. Although automation promises significant gains in speed, consistency, and analytical scale, capabilities strongly aligned with the Department of War's imperative to accelerate warfighting capability delivery, it also introduces new risks related to transparency, bias amplification, and accountability.

This paper conceptualizes this emerging dynamic as the "AI-to-AI procurement paradox." In such an environment, automated systems on both sides of the acquisition process may influence information inputs, proposal structures, and evaluation outcomes before human decision-makers can intervene. Without deliberate governance mechanisms, this dynamic could erode human reasoning, ethical oversight, accountability, and subject matter expertise that traditionally anchor federal procurement decisions.

---

**"AI IS RAPIDLY CHANGING OUR WORLD AND HAS SIGNIFICANT POTENTIAL TO TRANSFORM SOCIETY AND PEOPLE'S LIVES. HOWEVER, AI TECHNOLOGIES ALSO POSE RISKS THAT CAN NEGATIVELY IMPACT INDIVIDUALS, GROUPS, ORGANIZATIONS, COMMUNITIES, SOCIETY, AND ENVIRONMENTS." (GAO, 2025)**

---

### The Rise of AI in Federal Acquisition

Federal acquisition has historically adapted to technological advancements to improve efficiency and effectiveness, evolving from paper-based contracting to fully digitized procurement systems. AI represents the next significant evolution in this trajectory. Acquisition organizations are now experimenting with or deploying AI-enabled tools to analyze market intelligence, assist in drafting requirements, support pricing and cost realism assessments, and augment proposal evaluation processes.

These capabilities align closely with contemporary defense acquisition priorities, particularly the imperative to accelerate capability delivery in response to rapidly evolving threats. AI offers the ability to process large volumes of data at speeds unattainable by human analysts, identify trends across complex datasets, and promote consistency in evaluative decision-making. From an operational perspective, such efficiencies are attractive in an environment characterized by constrained resources, reductions in the size and experience of the acquisition and engineering workforces of the Department of War, and increasing demand for rapid acquisition outcomes.

However, the integration of AI into acquisition processes is also reshaping decision-making authority. As AI systems provide increasingly sophisticated recommendations, there is a growing risk that human decision-makers may defer to algorithmic outputs, especially under time pressure. This shift underscores the need to examine not only how AI is used, but also how it alters the balance between automation and human judgment in acquisition governance. iQuasar (2026) summarizes the underlying dilemma: "AI procurement in 2026 will transform federal contracting through automated evaluations, risk monitoring, and data-driven oversight." This anticipated transformation highlights the importance of understanding both the benefits and the challenges AI brings to acquisition governance, setting the stage for further discussion of its broader implications.

### The AI-to-AI Dilemma

The rapid adoption of AI within government acquisition is occurring alongside significant advancements in AI-enabled proposal development within industry. The defense industrial base is leveraging AI tools to automate compliance checks, generate technical narratives, optimize pricing strategies, and tailor proposals to solicitation requirements. These tools increase competitiveness and efficiency, but they also fundamentally change the nature of proposal



inputs entering the federal acquisition system.

The convergence of these trends creates an acquisition environment where automated government systems may increasingly assess, score, or filter proposals generated by automated industry systems. In this scenario, human involvement may be limited to overseeing machine-generated outputs rather than making substantive evaluative judgments. This dynamic raises several critical risks: transparency may be diminished if decision logic is embedded within opaque algorithms; algorithmic bias may be amplified if both proposal generation and evaluation rely on historically biased data; and accountability may be weakened, complicating responsibility for acquisition outcomes.

Most importantly, the AI-to-AI paradox threatens to erode the human judgment necessary for evaluating qualitative and context-sensitive factors such as innovation, operational risk, and mission alignment. These factors are central to defense acquisition success yet are difficult to fully capture through automated analysis. Without intentional safeguards and frameworks to navigate this dilemma, the pursuit of speed and efficiency may inadvertently undermine the integrity and trust that sustain effective acquisition.

### **Purpose and Thesis**

The purpose of this paper is to examine the AI-to-AI procurement paradox and assess its implications for federal acquisition integrity and warfighting capability delivery. By analyzing the parallel adoption of AI by government and industry, this research identifies the ethical, operational, and governance risks that emerge when automation increasingly mediates acquisition decisions. As noted, “Agencies still need a human in the loop ... but that person is reviewing something created instantly, instead of starting it from scratch manually or spending time filling in details in a blank template” (FNN, 2025). This shift underscores the importance of maintaining meaningful human oversight even as automation accelerates the acquisition process.

The central thesis of this paper is that while AI has the potential to significantly enhance acquisition speed and analytical rigor, it must not replace human judgment in critical decision-making processes. FAR 7.503(c)(12) explicitly classifies these responsibilities as inherently governmental, requiring direct human-in-the-loop involvement, and recent protest decisions reinforce that the source selection official must perform an independent assessment beyond the work of agency evaluators before making a final decision. Accordingly, AI should be deliberately integrated into a human-centered acquisition framework that preserves accountability, confirms ethical reasoning, and maintains sound strategic judgment. To address this challenge, this paper proposes a governance model called the Human Agency Assurance Framework (HAAF), designed to preserve meaningful human control within AI-enabled acquisition environments. The framework integrates oversight thresholds, verification and audit mechanisms, explainability requirements, governance structures, and workforce competencies to confirm that AI augments rather than replaces expert acquisition judgment.

### **Background: Guidance and Precedent**

The federal government has acknowledged the transformative potential of AI while recognizing the governance challenges it presents. Through various directives, frameworks, and guidance, it has sought to establish a foundation for responsible AI development and deployment. However, despite these efforts, risks associated with the use of AI for acquisition-specific functions, including procurement, integration, and life-cycle management, lack robust federal policy to regulate and optimize their use.

At the national level, America’s AI Action Plan underscores the importance of ensuring AI systems are ethical, transparent, and accountable (The White House, 2025). It highlights the





need for governance frameworks that address issues such as bias, privacy, and security. While the plan promotes innovation and public-private collaboration, it does not explicitly address the risks tied to the acquisition and operationalization of AI systems, particularly in complex environments like defense and national security.

The DOW has taken a more targeted approach to AI governance through its Ethical Principles for AI, which emphasize responsible, equitable, traceable, reliable, and governable AI (DOW, 2026). These principles are intended to guide the development and deployment of AI in defense applications. DOW's AI Strategy further outlines the need for robust testing, evaluation, and validation of AI systems to confirm they meet mission-critical requirements. However, gaps remain in addressing risks to the acquisition of AI systems, such as supply chain vulnerabilities, vendor accountability, and the long-term sustainment of AI technologies.

From an administrative perspective, the Office of Management and Budget (OMB) has issued guidance on the use of AI in federal agencies, including the OMB Memorandum M-21-06, which provides principles for AI use in government (GSA, n.d.). This guidance emphasizes transparency, accountability, and risk management but, like the documents referenced above, it does not delve into the unique challenges of acquiring AI systems, such as ensuring compliance with ethical standards during procurement or managing the risks of proprietary technologies.

Finally, Executive Order 14110 highlights the federal government's commitment to advancing responsible AI (GSA, 2025). It calls for interagency collaboration and the establishment of standards to govern AI use. However, like other federal directives, it does not specifically address the acquisition life cycle, leaving a critical gap in governance.

In summary, while federal directives, DOW frameworks, and administrative guidance collectively demonstrate a commitment to responsible AI governance, they do not fully address acquisition-specific risks. These risks include ensuring regulatory and ethical compliance during procurement, managing vendor relationships, mitigating supply chain vulnerabilities, and sustaining AI systems over time. This gap highlights the need for more comprehensive policies that integrate acquisition-specific considerations into the broader AI governance framework. Developing agency-specific Human Agency Assured Frameworks (HAAF) tailored to your organization is therefore a critical step, emphasizing the importance of human agency.

### **Precedent: Why Human Agency Matters in Acquisition**

Human agency is essential in the AI acquisition life cycle because it confirms that human judgment, oversight, and accountability remain central throughout procurement, deployment, and life-cycle management. In the context of acquisition, human agency is important for several key reasons, as shown in Table 1.



**Table 1. The Role Human Agency in Acquisition**

| Key Area                             | Description                                                                                                                                                                                                                                                                                          |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ethical Decision Making              | Human agency ensures that decisions about the selection, design, and deployment of AI systems align with ethical principles, such as fairness, accountability, and transparency.                                                                                                                     |
| Accountability and Responsibility    | Human agency ensures that individuals or organizations remain accountable for the outcomes of AI systems.                                                                                                                                                                                            |
| Risk Mitigation                      | The use of AI systems involves inherent risks, such as supply chain vulnerabilities, cybersecurity threats, and potential misuse of technology. Human agency is necessary to identify, assess, and mitigate these risks throughout the acquisition process.                                          |
| Contextual Understanding             | AI systems often lack the ability to fully understand the context in which they operate. Human institutional expertise bridges the gap between AI outputs and hallucinations.                                                                                                                        |
| Trust and Transparency               | Trust is a cornerstone of successful AI adoption, and human agency plays a vital role in building and maintaining that trust. When humans are actively involved in the acquisition process, they can ensure that AI systems are transparent, explainable, and aligned with stakeholder expectations. |
| Adaptability and Oversight           | AI systems are not static; they evolve over time through updates, retraining, or changes in operational environments. Human agency ensures that these changes are monitored and managed responsibly.                                                                                                 |
| Avoiding Over Reliance on Automation | Over reliance on AI systems without human oversight can lead to automation bias, where humans defer to AI decisions even when they are flawed.                                                                                                                                                       |

Scholars examining AI governance emphasize that maintaining human responsibility must begin at the earliest stages of system development and acquisition. Bode and Huelss (2025) state, “Retaining human responsibility in the development and use of AI-based systems requires establishing accountability mechanisms at the earliest stages of the lifecycle, ensuring that human agency is exercised across design, deployment, and oversight to prevent the diffusion of moral responsibility.” Human agency in acquisition is essential to confirm that AI systems are procured, deployed, and managed in alignment with ethical principles, risk mitigation, and accountability. It confirms that humans remain in control of critical decisions, fostering trust, transparency, and adaptability throughout the life cycle of AI systems. Without human agency, the acquisition process risks becoming overly reliant on automation, which can lead to ethical, operational, and security challenges. Next, this paper will take a deeper dive into the significant risks that arise in an AI-to-AI procurement ecosystem.

## Analysis: Risks of an AI-to-AI Procurement Ecosystem

**“GIVEN THE RAPID GROWTH IN CAPABILITIES AND WIDESPREAD ADOPTION OF AI, THE FEDERAL GOVERNMENT HAS TAKEN ACTION ... INCLUDING TO GUIDE AGENCIES IN ACCELERATING THE USE AND ACQUISITION OF AI.” (GAO, 2025)**

As federal acquisition organizations integrate AI into various stages of the acquisition life cycle, a set of interrelated risks emerges that extend beyond isolated technical concerns. These risks stem from how AI tools are embedded in the life cycle, governed across organizations, and relied upon by the workforce. In an AI-to-AI procurement ecosystem, these risks are compounded, resulting in systemic vulnerabilities that threaten transparency, competitive fairness, innovation, and alignment with the warfighting mission.

This section analyzes four primary risk categories: opaque reasoning and loss of transparency; algorithmic bias amplification; proposal homogenization and over optimization; and erosion of human judgment and contextual understanding. While analytically distinct, these



risks reinforce one another as AI adoption accelerates throughout all phases of the program and the entire acquisition life cycle.

### **Opaque Reasoning and Loss of Transparency**

A foundational principle of federal acquisition is the government's obligation to justify procurement decisions in a transparent, traceable, and defensible manner. Source selection outcomes must be defensible by contracting officers, oversight authorities, and, when necessary, judicial bodies. The introduction of AI-enabled tools into requirements development, evaluation, and decision support challenges this obligation when model outputs cannot be clearly explained or reproduced.

Many AI systems used for proposal scoring, risk management, or text generation operate as black boxes, with logic that is proprietary or opaque even to their developers (NIGP, 2026). These systems may produce precise rankings or narrative assessments without revealing which features were most influential or how trade-offs were resolved. When such outputs materially shape evaluations, acquisition officials may struggle to articulate why one offeror was rated superior, relying solely on the tool's conclusions. As Tillipman (2026) notes, "When agencies rely on AI systems whose conclusions cannot be meaningfully examined because the systems are opaque, technically complex, or contractually shielded, the government may be able to show what the tool produced, but not why those outputs reflect the reasoned judgment the APA requires." This opacity undermines legal defensibility under the Federal Acquisition Regulation (FAR), complicates debriefings, increases protest risk, and erodes confidence in the fairness of outcomes.

Transparency risks extend beyond evaluation to AI-assisted drafting of solicitations, requirements, and evaluation criteria. When AI tools generate or heavily influence acquisition documents, the government remains responsible for maintaining a defensible record of what inputs were used, which models were employed, how outputs were edited, and by whom. Weak logging, version control, or origin tracking complicates records management and post-award reviews. It also blurs the boundary between human judgment and machine recommendation, which is particularly problematic when ambiguities or errors in AI-generated language later become grounds for protest.

These challenges are magnified in high-consequence defense acquisitions. When AI-assisted analyses inform judgments about survivability, operational suitability, or system-of-systems integration, senior leaders and warfighters must understand the reasoning behind those assessments. The inability to clearly explain why a particular solution was deemed superior undermines confidence not only in the acquisition outcome, but also in the AI tools themselves.

### **Algorithmic Bias Amplification**

Algorithmic bias is a critical ethical and acquisition risk in an AI-to-AI procurement ecosystem, especially because bias can be amplified rather than mitigated when AI systems interact throughout the life cycle. In acquisition, AI models learn from historical data that reflects past award decisions, evaluation practices, and industrial base structures. This data is generated and stored across distributed government and industry systems, including SAM.gov, CPARS, agency contract writing systems, and program office repositories, vendor proposals, and market research artifacts. When these historical patterns are embedded in both government evaluation tools and industry proposal generation systems, bias becomes self-reinforcing.

On the government side, AI-enabled market research or proposal triage tools may implicitly favor vendors, approaches, or language resembling past winners, often privileging incumbents and established primes. On the industry side, proposal generation tools trained on



prior successful submissions optimize content to match those same patterns. When evaluative AI systems rely too heavily on similarity to past examples, rather than considering past performance as just one factor, innovative or unconventional solutions can be unfairly disadvantaged.

This feedback loop has direct implications for competitive fairness. Small businesses, the nontraditional defense industrial base, and other new entrants may find it increasingly difficult to compete if AI-mediated processes implicitly favor scale, legacy performance profiles, or standardized proposal structures. Conversely, traditional defense businesses, experienced small businesses, federally funded research and development centers (FFRDCs), UARCs, and universities could also be disadvantaged if automated processes are biased against their proposals and toward those of new entrants. Even when explicit demographic variables are excluded, proxies such as company size, contract history, or terminology can reintroduce disparate impacts. Over time, AI-to-AI interactions risk narrowing the competitive landscape, reinforcing incumbency advantages, and undermining statutory objectives related to fair and open competition, while giving the appearance of objective, data-driven decision-making. U.S. DoS (2024) notes that harmful bias and homogenization result in “amplification and exacerbation of historical, societal, and systemic biases; performance disparities between sub-groups or languages ... undesired homogeneity in data inputs and outputs resulting in degraded quality of outputs.” All of these factors complicate the overarching paradox that our proposed framework seeks to address.

### **Proposal Homogenization and Over Optimization**

The widespread adoption of AI in proposal development introduces strong incentives for over-optimization. Generative AI tools are designed to maximize compliance with solicitation requirements, align language precisely with evaluation criteria, and minimize perceived risk. While these capabilities can improve efficiency and baseline proposal quality, they also encourage convergence around a narrow set of “model-approved” narratives and structures.

As multiple offerors rely on similar tools trained on overlapping data from past solicitations and winning proposals, technical and management volumes increasingly resemble one another. MeriTalk (2026) highlights “a growing risk that AI could harm the integrity of the acquisition process, citing an uptick in protests and filings that rely on AI-generated content, including fabricated cases and citations—a byproduct of ongoing concerns about hallucinations.” Novel architectures, unconventional teaming strategies, or disruptive technologies may be normalized into familiar patterns or excluded entirely if they deviate from historical examples. When government evaluators also rely on AI-assisted tools optimized for pattern recognition, this homogenization is reinforced rather than challenged.

This dynamic poses a direct risk to innovation and long-term warfighting advantage. Kolt (2025) observes, “While the homogenization introduced by foundation models promotes efficiency ... it also gives rise to new risks familiar to complexity researchers.” Defense acquisition depends on identifying solutions that are not only compliant, but also operationally effective, adaptable, resilient, and responsive to evolving threats. An acquisition environment that implicitly rewards conformity and linguistic optimization over substantive differentiation risks selecting low-variance solutions that appear safe on paper but lack strategic depth. Over time, this suppresses experimentation and reduces the government’s exposure to transformative capabilities.

### **Loss of Human Judgment and Contextual Understanding**

The most significant risk in an AI-to-AI procurement ecosystem is the erosion of human judgment. Federal acquisition is inherently context-dependent, requiring experienced professionals to balance technical performance, cost, risk, schedule, and mission relevance.



However, AI tools introduce strong incentives for cognitive offloading, particularly under time pressure and resource constraints.

Automation bias can lead evaluators to over trust AI-generated scores, summaries, or narratives, especially when these are presented with quantitative precision or authoritative language. Over time, acquisition personnel may shift from exercising independent judgment to validating algorithmic outputs. This is particularly problematic for qualitative assessments such as innovation potential, management credibility, teaming viability, and adaptability to evolving operational needs, areas where context and experience are essential.

There is also a long-term institutional risk. If AI routinely drafts requirements, evaluation narratives, or acquisition strategies, newer professionals may never fully develop the tacit expertise required to detect unrealistic assumptions, hidden execution risks, or subtle misalignments with warfighter needs. AI systems trained primarily on historical data (e.g., prior RFPs and SOWs, source selection evaluation reports, CPARS past performance records, archived vendor proposals, etc.) cannot fully account for emerging concepts of operations, evolving adversary tactics, or cross-domain dependencies. Without deliberate human-in-the-loop engagement, acquisition decisions may optimize for what is easiest to measure rather than what is most critical to mission success. Additionally, transferring the institutional knowledge of subject matter experts with years of experience is extremely difficult, potentially leaving a significant gap in the AI's inputs.

This tension between analytical efficiency and human judgment is widely recognized in the broader AI governance literature. “The key challenge is to harness AI’s computational capabilities in crisis decision-making while preserving human moral agency. The goal is not to replace human decision-makers, but rather to create a collaborative interface where AI can support and enhance ethical reasoning without undermining fundamental moral responsibility” (Johnson, 2026). Over time, acquisition personnel may shift from exercising independent judgment to merely validating algorithmic outputs.

### **Compounding Effects Across the Acquisition Life Cycle**

Individually, each of these risks is manageable. Collectively, however, they form a compounding threat to acquisition integrity. Opaque reasoning obscures accountability; bias amplification distorts competition; proposal homogenization suppresses innovation; and diminished human judgment weakens mission alignment. As AI is applied sequentially, from requirements framing, to proposal generation, to evaluation, these dynamics reinforce one another, creating systemic vulnerabilities that are difficult to detect and correct after the fact.

Addressing the AI-to-AI procurement paradox therefore requires more than technical safeguards applied at isolated points in the process. As the Open Contracting Partnership (2025) states, poor AI use in procurement “can cost governments millions, not just in failed projects, but in lost public trust.” Addressing these risks demands a holistic, life-cycle-oriented approach that preserves transparency, fairness, and human agency as AI becomes embedded in federal acquisition. The following section builds on this analysis by proposing a framework for maintaining meaningful human oversight within AI-augmented procurement environments.

### **Proposed Framework and Recommendations: Preserving Human Agency in AI-Enabled Acquisition**

Federal acquisition organizations are rapidly integrating AI into all aspects of the acquisition life cycle. While these applications can reduce lead times and enhance analytical capabilities, they also introduce new failure modes. Recent federal direction emphasizes risk management, transparency, and accountability practices for AI, especially for uses that impact public rights and safety (OMB, 2024).





### Proposed Human Agency Assurance Framework (HAAF)

The proposed HAAF pairs acquisition decision authority with measurable AI assurance gates. It is designed to be implemented within existing acquisition governance structures (e.g., source selection plans, quality assurance surveillance plans, and risk registers) and aligns with federal AI governance expectations for inventorying AI use cases, managing risks, and documenting controls (OMB, 2024). Figure 1 illustrates the core principles of the proposed HAAF from a source selection perspective.




*Figure. 16 HAAF for Source Selection*

HAAF consists of five integrated components, as shown in Table 2.

**Table 2. Components of the HAAF**

| <b>Component</b>            | <b>Mechanism</b>                                                                                                                                                                  | <b>Purpose</b>                                                                                                                                                                      |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Oversight Thresholds</b> | Human-in-the-Loop (HITL) and Human-on-the-Loop (HOTL) oversight tiers that define when humans must approve decisions, supervise automated processes, or directly execute actions. | Confirms that the level of human control scales with operational risk, mission impact, and ethical considerations, preserving meaningful human judgment in consequential decisions. |
| <b>Assurance Gates</b>      | Verification, Validation, and Accreditation (VV&A) processes are combined with audit and traceability requirements throughout system development and deployment.                  | Provides evidence that AI systems function as intended, maintain transparency in decision logic, and meet reliability, safety, and compliance standards prior to operational use.   |
| <b>Governance</b>           | Provides evidence that AI systems function as intended, maintain transparency in decision logic, and meet reliability, safety, and compliance standards prior to operational use. | Establishes institutional controls to prevent misuse, confirm responsible deployment, and maintain alignment with organizational objectives and ethical standards.                  |
| <b>Policy and Guidance</b>  | Formal AI policies, acquisition guidance, and implementation frameworks that translate strategic objectives into operational rules and procedures.                                | Provides consistent standards for integrating AI into acquisition and operational environments, ensuring compliance with legal, ethical, and regulatory requirements.               |
| <b>Workforce Readiness</b>  | Role-based AI literacy training, competencies, and operating procedures for acquisition professionals, engineers, and oversight personnel.                                        | Sustains accountability and informed decision-making by equipping the workforce with the knowledge required to evaluate, supervise, and govern AI-enabled systems.                  |

Before examining the five components, it is essential to emphasize that the success of HAAF depends on data quality. This requires a clear understanding of the intended tool, the use of relevant and high-quality data, and sound model development. Aligning data with desired outcomes from the outset confirms the model is effective before training even begins. With that foundation, the five HAAF components are outlined below.

### **Human-in-the-Loop and Human-on-the-Loop Thresholds**

Preserving human agency begins by defining which acquisition decisions can be AI-assisted versus AI-automated and, critically, where humans must remain the accountable decision-makers. In acquisition, “human agency” means that contracting officers and designated acquisition officials retain meaningful control over outcomes, can understand and challenge AI supported recommendations, and can intervene before harm occurs. Practical implementation requires explicit thresholds tied to decision criticality, uncertainty, and potential harm. HAAF uses a four-tier oversight model, as shown in Table 3.



**Table 3. Tiered Tool Levels**

| Tier          | Oversight             |
|---------------|-----------------------|
| <b>Tier 0</b> | Automation Prohibited |
| <b>Tier 1</b> | Human-in-the-Loop     |
| <b>Tier 2</b> | Human-on-the-Loop     |
| <b>Tier 3</b> | Administrative AI     |

**Tier 0 (Automation prohibited):** AI may support drafting or retrieval but cannot recommend or execute decisions (e.g., determinations affecting eligibility, responsibility, or human rights and safety). This tier is appropriate for high-stakes, high-impact use cases and mirrors the federal emphasis on heightened practices for uses that affect human rights and safety (OMB, 2024).

**Tier 1 (HITL approval required):** AI may generate recommendations (e.g., draft evaluation narratives, risk assessments, price realism indicators), but a human must review the rationale and explicitly approve the output before it is used. HITL is required when the decision has material consequence, when data quality is uncertain, or when the model is newly deployed or recently updated.

**Tier 2 (HOTL supervision required):** AI may execute bounded tasks with continuous monitoring and rapid intervention capability (e.g., automated compliance checks against solicitation instructions; anomaly detection that flags transactions). HOTL is appropriate when the task is repetitive and low-consequence per instance, but aggregate errors could create material risk.

**Tier 3 (Human-out-of-the-loop (HOOTL) allowed only for non-material administrative support):** AI may operate with minimal oversight where errors are easy to detect and correct and have negligible impact (e.g., formatting, routing, de-duplication of non-decisional records).

Operationally, each AI-enabled acquisition should be assigned a tier by reviewing the following questions:

- Does the output affect award decisions, contractor selection, performance ratings, payment integrity, or rights/safety?
- Does the use of the AI systems significantly increase the likelihood of a protest of the award?
- If incorrect, can the output be corrected before harm occurs and without major cost/schedule/mission impact?
- Is the model operating outside the bounds of validated data (e.g., new commodity, new mission area, novel solicitation structure)?
- Can a reviewer replicate or meaningfully discern the basis for the recommendation?
- Is the acquisition exposed to intentional manipulation (e.g., prompt injection, data poisoning, strategic responses)?



There are two clear “tripwires” that require HITL oversight: (1) any use in source selection judgments (e.g., technical acceptability, best-value trade-offs, limiting eligible sources, past performance narratives, sole source justifications), and (2) any use that could influence enforcement or eligibility decisions. These controls complement federal accountability practices and reflect the acquisition community’s consensus that AI should augment, not replace, expert judgment in core decisions (MITRE, 2024). Acquisition practitioners increasingly describe AI as a decision support capability rather than a replacement for contracting authority. As Hubbard (2025) notes, “AI serves as a decision support tool, not a decision-maker. Acquisition professionals retain full authority over all procurement decisions, with AI providing analysis and recommendations that humans validate and approve.”

This can be achieved by requiring an AI Oversight Annex in acquisition and source selection plans that lists AI-enabled aspects, assigns the oversight tier, identifies accountable officials, and defines escalation triggers (e.g., confidence below threshold, drift indicators, team findings). This aligns with government-wide expectations for AI use case inventories and risk management documentation (OMB, 2024).

### **Verification, Validation, and Audit Requirements**

To preserve agency, humans must be able to trust AI outputs but also verify them. VV&A should be scaled according to the oversight tier and associated risk. NIST’s Generative AI Profile for the AI Risk Management Framework outlines governance and measurement activities that help organizations identify and manage generative AI risks through repeatable evaluations (NIST, 2024). Similarly, OMB directs agencies to implement minimum risk management practices and documentation for AI uses, with heightened expectations for applications that impact human rights or safety (OMB, 2024).

The HAAF establishes four assurance gates that must be satisfied before an AI capability is used:

**Gate A – Intended use definition and data pedigree:** Document purpose, decision boundaries, prohibited uses, and data provenance. Include training data sources (or vendor attestations), sufficiency of available data, data rights, privacy constraints, and known limitations.

**Gate B – Pre-deployment validation:** Validate accuracy and robustness against representative acquisition artifacts (e.g., solicitations, proposals, market research summaries) and stress test for manipulation (e.g., prompt injection, instruction-following attacks). For scoring or ranking use cases, validate calibration and evaluate differential error rates across vendor types where feasible.

**Gate C – Operational monitoring and drift management:** Monitor performance over time (e.g., concept drift, data drift), track user overrides, and establish rollback criteria. Where AI influences decisions, log the “reason codes,” context, and evidence presented to the reviewer.

**Gate D – Independent audit and accountability:** Periodically conduct independent reviews of model performance, security controls, and documented decisions. GAO’s AI Accountability Framework emphasizes practices around governance, data, performance, and monitoring, and can serve as an audit checklist for federal AI deployments (GAO, 2024). GAO has also highlighted the rapid growth of generative AI use at agencies and the need for policy-compliant management amid fast technical change (GAO, 2025).

Auditability is particularly important for acquisition artifacts that may be sensitive to protests. The HAAF recommends that AI tools used in evaluation support produce defensible records, including versioned outputs, citations to underlying inputs, reviewer decisions (accept/modify/reject), and explanations of how the AI was used. This does not require full



model transparency for proprietary systems; rather, it requires decision traceability sufficient for oversight and after-action review. Legal scholars have similarly warned that agencies must preserve the ability to interrogate and override AI tools used in procurement decisions. As Tillipman (2026) notes, “Agencies that deploy AI tools should therefore secure contractual rights that preserve the conditions for independent judgment: sufficient information to understand how the tool applies the solicitation’s evaluation criteria and the practice ability to question, and where necessary, override or disregard its outputs.”

### ***Require Explainability and Decision Traceability***

AI-assisted evaluations must be transparent and interpretable. Explainability artifacts, including rationale summaries, confidence indicators, or input-output mappings, should accompany AI outputs and be retained in the administrative record. Transparency and interpretability are core elements of responsible AI governance (NIST, 2023). In acquisition contexts, explainability allows evaluators to demonstrate how AI informed decisions while ensuring that humans retain authority.

Explainability also strengthens protest defensibility by documenting the reasoning behind evaluation judgments. Without traceability, agencies risk relying on impervious outputs that cannot be justified under scrutiny.

### **AI Governance and Oversight**

AI governance in acquisition should be treated as a system of controls encompassing people, process, and technology. OMB directs agencies to strengthen AI governance by establishing internal capabilities and implementing minimum risk management practices for certain AI uses (OMB, 2024). Complementary assurance guidance emphasizes continuous monitoring, documentation, and transparency regarding context and limitations to sustain accountability over time (NIST, 2024). In acquisition, governance must also preserve the integrity of the procurement record, ensuring that AI assistance does not introduce undocumented criteria, untraceable reasoning, or automation bias. Moreover, as demonstrated in the development of complex technical regulatory regimes, the best outcomes result from continuous and substantive dialogue among tool developers, acquisition personnel, and regulatory bodies.

AI governance includes ethical safeguards, which should be expressed as enforceable requirements rather than aspirational principles. The EU Artificial Intelligence Act (AI Act) adopts a risk-based approach with obligations for high-risk systems, offering a concrete template for procurement-facing controls (European Parliament and Council, 2024). For federal agencies, OMB’s direction similarly stresses risk management practices and governance structures for AI uses that may affect rights and safety (OMB, 2024). These considerations should be incorporated when formulating the governance and oversight associated with AI usage. The recommendations below are designed to be implementable through familiar acquisition governance artifacts and to scale across multiple AI use cases.

### ***Establishing Cross-Functional AI Oversight Boards***

Create a cross-functional AI Oversight Board that integrates acquisition, program management, legal, privacy, cybersecurity, data governance, and independent testing and audit functions. The board’s role is to set and maintain guardrails that confirm AI use is repeatable and defensible. At a minimum, the board should: maintain an inventory of AI-enabled acquisition use cases; assign oversight tiers and approval authorities by risk; approve validation plans and monitoring metrics; consult with industry on use of AI in acquisition processes; maintain awareness of state-of-the-practice AI capabilities applicable to acquisition activities; and review incidents, drift indicators, and material user overrides. This structure supports OMB’s emphasis on agency AI governance and aligns with accountability practices recommended by GAO for





federal AI use. To further enhance its impact, the board should publish a brief quarterly “AI acquisition governance bulletin” describing approved use cases, controls, and lessons learned to create institutional learning loops (GAO, 2025).

### ***Embedding Governance in Contract Terms***

When agencies procure AI tools or AI-enabled services, governance must be translated into enforceable contract language. Contracts should require contractors to support verification and monitoring activities, provide audit evidence, and notify the government of significant model changes. The government should use a concrete, risk-based set of obligations, including documentation, logging, human oversight, and robustness, for higher-risk systems to inform procurement requirements.

Minimum governance contract terms for AI acquisitions should include model and prompt configuration control and version disclosure, logging requirements and data retention for audit purposes, performance monitoring and drift notification, security testing for adversarial threats, and an obligation to support independent government assessments. These terms operationalize accountability and improve traceability of AI-assisted work products. Governance may also include a “model change clause” requiring advance notice and government approval for changes that could materially affect outputs used in acquisition decisions (e.g., evaluation support, fraud indicators, compliance screening).

### ***Defining Qualitative Judgment Anchors***

Acquisition activities often rely on qualitative judgments, such as best-value trade-offs, relevance of past performance, experience with previous acquisition decisions and their programmatic outcomes, or risk assessments. To preserve human agency, agencies should define qualitative judgment anchors—short, plain-language criteria that clarify what “good” looks like for common qualitative determinations and specify the evidence required to support them. Anchors reduce automation bias by requiring reviewers to map AI outputs to meaningful criteria, rather than accepting fluent text as a proxy for correctness.

Anchors can be incorporated into evaluation worksheets and Standard Operating Procedures (SOPs). For example, a past-performance narrative anchor might require the reviewer to cite the reference, provide a relevance rationale aligned with scope and complexity, and include corroborating record specifics. NIST’s Generative AI profile emphasizes the importance of context, limitations, and measurement in managing generative risks (NIST, 2024). Anchors deliver this contextualization at the decision point. Additionally, the government could require an “AI-assisted work product checklist” for any AI-generated narrative used in the procurement record, confirming citations, factor alignment, and exclusion of non-evaluated criteria (GAO, 2025).

### ***Requiring Industry to Disclose AI Use and Metadata***

Human agency and fairness are compromised when the government cannot distinguish between vendor-authored content and AI-generated content, or when AI is used to optimize submissions in ways that conceal weaknesses or introduce unverifiable claims. Agencies should require offerors and contractors to disclose material AI use and associated metadata for acquisition-relevant artifacts (e.g., proposals, technical volumes, deliverable narratives, and automated test/analysis outputs).

A practical disclosure approach is to scope the requirement to material AI use, defined as either substantive content generation or automated analysis, and require disclosure of the tools used and their purpose; the tool and model version; the date, time, and high-level configuration or policy mode, if relevant; and a contractor attestation that all factual claims were human-verified. For high-dollar or higher-risk procurements, requiring an “AI content



provenance annex” that identifies AI-generated sections and documents human verification steps performed by the offeror would be ideal.

## **AI Policy and Guidance**

Governance mechanisms are most effective when reinforced by clear policy and repeatable guidance. GAO (2025) has observed that agencies face challenges complying with evolving federal AI policies while keeping pace with rapid technical change. Developing practical policy instruments, external guidance for industry, and internal acquisition templates and SOPs reduces inconsistency and improves audit readiness.

### ***Issuing Formal AI Policy Guidance to Industry***

Agencies should issue formal guidance to industry clarifying expectations for acceptable AI use in proposal preparation and performance, required disclosures and metadata, security and privacy constraints for contractor tools, and how AI-related claims will be evaluated. This guidance can be delivered through acquisition policy memoranda, standardized solicitation attachments, and industry day materials. It should include concrete examples of what constitutes “material AI use,” what evidence is expected for AI-produced analyses, and what logs or artifacts must be retained for post-award audit. Doing so reduces ambiguity, levels the playing field, and supports OMB’s push for agency AI governance and risk management (OMB, 2024).

External organizations have also recommended shared guidance mechanisms that streamline safe AI procurement across government, such as a national guidance platform for AI acquisition (RAND, 2025). Agencies can adopt this approach by publishing a lightweight “AI Acquisition Guide for Vendors” and updating it regularly. In addition, the government could include a standard “AI Instructions to Offerors” attachment for AI-relevant procurements that covers disclosure, data handling, security, and evaluation evidence requirements (NIST, 2024).

### ***Developing Internal Federal AI Acquisition Policies, Templates, and SOPs***

Internally, agencies should build a reusable AI acquisition kit—templates, clauses, evaluation factors, checklists, and SOPs—that aligns with government and agency-wide direction and scales across contracting activities. The kit should include an “AI Oversight Annex” template for acquisition plans and source selection plans; standard language for VV&A evidence, monitoring, and audits; a QASP addendum for AI-enabled contractor workflows and deliverables; and incident response playbooks for AI failures.

OMB’s M-24-10 memorandum provides a baseline for governance and risk management practices that should be reflected in internal procedures (OMB, 2024). NIST’s Generative AI profile offers actionable risk categories and control ideas that can be translated into SOP checklists and acceptance criteria (NIST, 2024). GAO’s accountability practices can serve as an audit-aligned backbone for these templates, ensuring traceability and continuous monitoring (GAO, 2024). GAO’s recent work also highlights how data quality and workforce capability determine whether AI can be leveraged responsibly in oversight-heavy domains, lessons that apply directly to acquisition integrity and payment controls (GAO, 2026). Finally, agencies should implement a “tiered authorization” model in which staff may use low-risk AI tools by default, but acquisition use cases with material decision influence require documented approval and training authorization.

## **AI Literacy for the Acquisition Workforce**

Human agency fails in practice when the workforce cannot recognize AI limitations, challenge outputs, or operate monitoring controls. AI literacy for acquisition is not a generic introduction to AI; it is role-based competence tied to acquisition decisions. GAO has noted that the rapid expansion of generative AI use across agencies creates management challenges and



underscores the need for policy-compliant adoption and workforce readiness (GAO, 2025). DAU reports similar findings, highlighting practical GenAI use cases in acquisition and the importance of training to use these tools effectively (DAU, 2025). HAAF recommends a role-based literacy model aligned to acquisition roles:

- **All hands (baseline):** understanding what AI can or cannot do; data sensitivity; hallucination risk; prompt hygiene; and mandatory disclosure rules.
- **Practitioners (1102s, CORs, analysts):** interpreting AI outputs; detecting red flags; validating citations; understanding uncertainty; and documenting AI use in acquisition records.
- **Supervisors and Source Selection Authority (SSA) support staff:** selecting oversight tiers; setting acceptance criteria; approving use cases; and managing risk registers.
- **Technical and governance leads (CAIO/IT/security):** model governance; monitoring; security testing; incident response; and audit readiness.

A practical literacy requirement is the ability to run “sanity checks” and structured reviews. For example, when AI drafts evaluation narratives, reviewers should confirm that every factual claim traces to the proposal or record, the narrative aligns with stated evaluation factors, and the tool does not introduce prohibited comparisons or non-evaluated criteria. When AI supports market research, reviewers should check source diversity, recency, and whether the tool’s summary omits disconfirming evidence. Finally, literacy must include organizational learning loops. MITRE (2024) argues that acquisition should blend the “art” of expert judgment with AI’s ability to reduce tedious tasks, provided organizations proceed cautiously with testing and integration. Institutionalizing after-action reviews, override analysis, and periodic refresher training helps sustain this cautious adoption posture.

Every agency should establish mandatory, role-based AI training tied to authorization for using Tier 1 or Tier 2 tools; incorporate AI use and oversight procedures into SOPs and Quality Assurance Surveillance Plans (QASPs); and conduct periodic exercises, such as tabletops or red teams, to practice intervention and documentation under realistic acquisition scenarios (DAU, 2025).

AI can improve acquisition speed and analytical capacity, but only if governance preserves accountable human judgment. The HAAF provides an actionable approach by defining oversight tiers for each use case, requiring VV&A evidence proportional to risk, and equipping the workforce with the knowledge to challenge and supervise AI systems. These steps align with current federal governance direction and emerging assurance guidance, while protecting the integrity and defensibility of acquisition outcomes.

## Operationalizing the Framework: Use Cases

Conceptual frameworks for responsible AI adoption in federal acquisition must ultimately be validated through practical application. While the preceding sections identified risks inherent in an AI-to-AI procurement ecosystem, acquisition leaders require concrete examples demonstrating how AI can be integrated into real-world workflows without displacing human judgment or undermining accountability. This section operationalizes the proposed framework through applied use cases that illustrate how narrowly scoped AI capabilities, implemented as micro-tools, can enhance acquisition effectiveness while preserving transparency, traceability, and human agency.



## Veterans Affairs Example

The Department of Veterans Affairs (VA) provides a compelling hypothetical example of how AI can be responsibly integrated into early-phase acquisition activities while strictly preserving human agency. In a recent market research effort, VA's acquisition team issued a Request for Information (RFI) to gather industry insights on a complex, mission-critical capability, supporting Veteran healthcare delivery through advanced technological solutions (e.g., interoperability, data analytics, or clinical workflow enhancements aligned with VA's digital modernization goals). The RFI elicited over 75 detailed vendor responses, each averaging 50 pages of technical narratives, capability statements, past performance examples, and innovative concepts. Manually reviewing and synthesizing this volume of unstructured, heterogeneous material would have required weeks of effort from a small team of contracting specialists, program managers, and SMEs, risking fatigue-induced oversights and schedule slippage in an environment where rapid requirements refinement is essential to accelerating Veteran capabilities.

To address this challenge without ceding evaluative control, VA deployed a narrowly scoped AI micro-tool, a generative AI-assisted analysis pipeline configured as a decision-support utility rather than an autonomous evaluator. The tool, built on secure LLM infrastructure with strict data handling protocols, performed the following bounded functions:

- **Key Data Extraction:** The AI parsed each RFI response to automatically extract structured elements such as proposed technical approaches, compliance with VA technical reference models, estimated maturity levels (e.g., TRL/MRL), teaming structures, cybersecurity postures, and rough-order-of-magnitude (ROM) cost ranges.
- **Theming and Clustering:** Using natural language processing techniques, the tool identified emergent themes across responses (e.g., common barriers to interoperability, recurring emphasis on cloud-native architectures, concerns about legacy system integration, or innovative uses of predictive analytics for Veteran outcomes). It grouped similar concepts, quantified theme prevalence (e.g., "65% of respondents highlighted FHIR-based standards"), and surfaced outliers (e.g., novel, non-traditional approaches from small businesses).
- **Summary Reporting:** The AI generated concise, evidence-based synopses for each response, with hyperlinks to original document sections, page numbers, and quotes. Importantly, the micro-tool operated under explicit HITL constraints consistent with the proposed HAAF Tier 1 oversight: AI outputs were never used directly in decision-making. All generated summaries, themes, and extractions were treated as preliminary working aids requiring mandatory human review and validation.

The acquisition team, consisting of 1102s, CORs, SMEs, and clinical/technical leads, then took ownership of the process through a structured, multi-stage workflow:

1. **Initial Validation Sprint:** In collaborative sessions, evaluators cross-checked a statistically sampled subset (approximately 20%) of AI-extracted data against source documents to confirm accuracy, detect any hallucinations or misinterpretations (e.g., conflating vendor marketing language with verifiable capabilities), and refine prompt parameters for subsequent runs. This calibration step achieved over 95% alignment after one iteration.
2. **Thematic Synthesis and Gap Analysis:** Human reviewers examined the AI-generated theme clusters to draw contextual conclusions. For instance, while the tool highlighted widespread industry support for API-first designs, VA experts identified that many responses underestimated VA-specific data sovereignty and Veteran privacy



requirements under 38 U.S.C. and HIPAA. This human insight led to strengthened requirements language around federated identity and zero-trust principles.

3. **Impact on Requirements Development:** The combined AI-assisted theming and human evaluation directly informed refinement of the draft Performance Work Statement (PWS) and evaluation criteria. Feedback revealed industry concerns about VA's legacy integration challenges, prompting adjustments to emphasize modular, open-architecture solutions and phased deployment incentives. It also surfaced opportunities for non-traditional entrants, leading to targeted small business set-aside considerations and outreach to underrepresented vendors.
4. **Documentation and Traceability:** Every final decision, whether accepting, modifying, or rejecting an AI-flagged theme, was logged with rationale, linked back to primary source material, and incorporated into the acquisition plan's AI Oversight Annex. This confirmed full defensibility during internal reviews, potential protests, or GAO scrutiny.

The outcome was transformative yet fully human-centered: market research that traditionally consumed 8–12 weeks was compressed to approximately 2–3 weeks, freeing capacity for deeper industry engagement (e.g., one-on-one vendor dialogues and solution demonstrations). More critically, the AI micro-tool acted purely as an accelerator for pattern recognition and data reduction, never as a substitute for judgment. Human experts retained authority to interpret qualitative nuances (e.g., assessing the credibility of a vendor's claimed innovation versus hype), weighed mission alignment against Veteran outcomes, and made binding determinations on requirements direction.

This VA use case exemplifies HAAF principles in practice: Tier 1 HITL for material market research influencing downstream source selection; scaled VV&A through sampling, calibration, and audit logging; and role-based AI literacy enabling the workforce to challenge outputs effectively. By treating AI as a force multiplier for human insight rather than a decision proxy, VA preserved accountability, mitigated risks of homogenization or bias amplification, and confirmed that industry feedback meaningfully shaped an acquisition approach better aligned with accelerating high-impact capabilities for Veterans. This disciplined integration demonstrates that responsible AI adoption can enhance speed and analytical depth in federal acquisition without eroding the human command essential to ethical, mission-successful outcomes.

### **FFRDC Example**

A central reference point for the HAAF is the Acquisition AI Research (ACQAIRE) initiative developed by the MITRE Corporation. ACQAIRE provides a research and experimentation environment focused specifically on acquisition-relevant AI use cases, enabling government, industry, and FFRDC stakeholders to explore how AI can support acquisition tasks under disciplined governance and human-in-the-loop conditions.

### ***Micro-Tools as Decision Support, Not Decision Replacement***

A foundational principle of the proposed framework is that AI should be deployed as a collection of narrowly scoped, task-specific micro-tools rather than as an end-to-end decision-making system. ACQAIRE exemplifies this approach by emphasizing AI capabilities that assist acquisition professionals in performing discrete functions, such as document triage, evidence identification, or cross-volume consistency analysis, without automating evaluative judgments.

Within an ACQAIRE-like environment, AI tools are deliberately constrained to surface information rather than render conclusions. For example, instead of generating composite scores or ranking offerors, AI may identify passages relevant to specific evaluation factors or flag potential areas of interest for further human review. Evaluators retain full authority to interpret the evidence, weigh trade-offs, and determine relevance to mission needs. This design





choice directly mitigates automation bias and reinforces the role of acquisition professionals as the final arbiters of value and risk.

This micro-tool paradigm is transferable across acquisition contexts. Government acquisition offices can use such tools to manage evaluator workload during complex source selections; industry teams can employ similar capabilities internally to improve proposal coherence and compliance without overstating claims; and FFRDCs can serve as neutral venues for testing and validating tool behavior before broader adoption. In each case, AI augments cognition and efficiency without supplanting judgment.

### ***Evidence Tagging and Traceability in Evaluation Workflows***

One of the most significant contributions of ACQAIRE-aligned use cases is the emphasis on evidence tagging and traceability. A persistent challenge in AI-assisted acquisition is ensuring that outputs remain directly linked to authoritative source material. ACQAIRE addresses this challenge by associating AI-identified observations, such as potential strengths, weaknesses, or risks, with specific proposal text, page locations, and evaluation criteria.

This approach enhances transparency and legal defensibility. Evaluators can review AI-suggested annotations, accept or reject them, and document their reasoning while maintaining a clear audit trail that demonstrates how conclusions were reached. During debriefings, protests, or oversight reviews, acquisition officials can point to documented evidence chains rather than opaque model outputs. AI thus becomes a tool for strengthening documentation rigor rather than obscuring decision logic.

Evidence tagging also reduces the risk of compounding errors throughout the life cycle. By ensuring that summaries and analytical outputs remain anchored to original proposal content, ACQAIRE-style tools mitigate the effects of hallucination, over-compression, and mischaracterization. Evaluators are consistently prompted to engage with primary source material, reinforcing human oversight and contextual understanding.

### ***Evaluator Support and Cognitive Load Reduction***

Complex source selections place substantial cognitive demands on evaluation teams, particularly in time-constrained acquisition environments. ACQAIRE use cases demonstrate how AI micro-tools can reduce cognitive load while preserving evaluative rigor. Examples include tools that highlight inconsistencies across proposal volumes, flag potential disconnects between technical and management approaches, or identify deviations from solicitation requirements.

Importantly, these tools do not determine compliance or merit. Instead, they function as structured prompts that help evaluators focus their attention where it is most needed. By reducing time spent on mechanical cross-checking, AI enables acquisition professionals to concentrate on higher-order judgments involving innovation, risk, and mission alignment. This supports acquisition speed without sacrificing the qualitative assessments essential to warfighting effectiveness.

From a workforce development perspective, this approach also preserves and reinforces expertise. Junior evaluators gain insight into how experienced professionals assess evidence and trade-offs, while senior evaluators remain responsible for final judgments. AI's role is explicit and bounded, supporting learning rather than eroding skill.

### ***Applicability Across Government, Industry, and FFRDC Environments***

The ACQAIRE model demonstrates that the proposed framework is feasible across diverse acquisition environments. In government acquisition offices, AI micro-tools can be integrated into existing evaluation processes with minimal disruption, provided that governance



mechanisms for logging, version control, and human-in-the-loop decision points are established. In industry settings, similar tools can be applied during internal reviews to identify unsupported claims, inconsistencies, or compliance gaps before proposal submission, reducing downstream execution and accountability risks.

FFRDCs such as MITRE play a distinctive role in this ecosystem. Positioned at the intersection of policy, technology, and acquisition practice, they provide neutral environments for experimentation, red-teaming, and validation. ACQAIRE illustrates how FFRDC-led research efforts can translate high-level AI policy principles into operationally viable tools, offering acquisition organizations a lower-risk pathway to responsible AI adoption.

### ***Implications for Framework Adoption***

The ACQAIRE use cases demonstrate that the AI-to-AI procurement paradox is not an unavoidable consequence of automation, but rather the result of design and governance choices. When AI is implemented as a set of constrained, transparent, and traceable micro-tools, it can enhance evaluation quality and acquisition speed while preserving accountability and human agency.

By grounding the conceptual framework in real-world practice, ACQAIRE shows that responsible AI adoption in federal acquisition is achievable using existing technologies. The key is disciplined integration: clearly defined human-in-the-loop thresholds, robust evidence traceability, and deliberate limitation of AI authority, as outlined by the proposed HAAF.

### **Recommendations and Policy Implications**

**Establish a Federal AI Acquisition “Circuit Breaker” Protocol:** The HAAF prescribes oversight tiers but lacks a mechanism for addressing AI system failures during production in mid-acquisition. A circuit breaker policy would require agencies to predefine automatic suspension triggers, such as a confidence score dropping below a validated threshold, a drift indicator activating, or a protest citing AI-generated content. These triggers would immediately pause AI-assisted evaluation and revert to manual processes. The key insight is that the HAAF’s assurance gates are pre-deployment controls; this protocol fills the gap for real-time failure response during active source selections, where stopping and restarting can have significant cost and schedule consequences.

**Establish an Independent Federal AI Acquisition Auditor Function:** The HAAF’s Gate D calls for independent audits but does not assign any institution the authority or resources to conduct them. An independent Federal AI Acquisition Auditor, potentially within GAO, would have standing authority to audit AI tool deployments in federal acquisition, review protest records for AI-related patterns, and publish public findings. The critical design principle is that this function must be independent of the agencies being audited, which the HAAF’s current internal board structure does not guarantee. This mirrors the logic behind the Inspector General system: internal governance is necessary but not sufficient when public accountability is at stake.

**Integrate HAAF Tiers into DODI 5000.02:** The Under Secretary of War for Acquisition and Sustainment (USW(A&S)) should amend DOD Instruction 5000.02 to require HAAF-tier assignments across MDAPs, OTAs, and rapid acquisition pathways. This memo establishes policy and prescribes procedures for managing acquisition programs from initiation through sustainment (DODI 5000.02, 2020). While DODI 5000.02 defines milestone decision authorities, documentation requirements, and oversight thresholds for programs, it does not address how AI tools used within those acquisition processes should be governed.

**Establish a Sunset and Revalidation Requirement for Deployed AI Acquisition Tools:** The HAAF addresses pre-deployment validation and ongoing monitoring but lacks a mechanism for



periodic revalidation. AI models degrade, training becomes outdated, and acquisition environments evolve; a tool validated in 2024 may behave very differently when applied to 2027 proposals, potentially without any visible failure signal. A mandatory sunset provision would require every AI tool used in acquisition to undergo full revalidation against current data on a defined cycle, with automatic suspension if revalidation is not completed. This closes what is perhaps the most institutionally underappreciated gap in the HAAF: the assumption that a tool that passed assurance gates at deployment remains trustworthy indefinitely. To date, this has not been tested in a protest or court venue, and new limits on the use of AI may be imposed later depending on such outcomes.

## Conclusion: Keeping Humans in Command

AI will continue to transform federal acquisition by dramatically expanding the government's ability to process information, evaluate alternatives, and accelerate procurement timelines. However, integrating AI into acquisition workflows introduces governance challenges that cannot be addressed through technology alone. As automated tools increasingly influence both proposal generation and evaluation, maintaining meaningful human agency becomes essential to preserving the integrity, transparency, and accountability of federal procurement. "The future of federal procurement is about achieving mastery in both AI efficiency and human judgment. The winners will be those organizations that act now" (Shipley Associates, 2026).

This paper introduced the concept of the AI-to-AI procurement paradox to describe the emerging interaction between government and industry AI systems within acquisition processes. If left unmanaged, this risks amplifying bias, obscuring evaluative reasoning, and eroding the expert judgment required for mission-aligned procurement decisions. To mitigate these risks, the HAAF provides a structured governance approach for integrating AI into acquisition while preserving human control. By establishing explicit oversight tiers, implementing verification and audit mechanisms, strengthening explainability and traceability requirements, embedding governance expectations into policy and contracts, and developing an AI-literate acquisition workforce, agencies can responsibly harness AI capabilities without compromising procurement integrity.

Ultimately, the objective is not to slow the adoption of AI within federal acquisition, but to guide it. The warfighter's need for speed must be balanced with institutional safeguards that confirm acquisition decisions remain defensible, ethical, and strategically sound. When implemented responsibly, AI can serve as a force multiplier for human expertise rather than a substitute for it. Keeping humans in command confirms that technological acceleration strengthens, rather than undermines, the trust and accountability upon which effective defense acquisition depends.

## References

- Ada Lovelace Institute. (2024, October 1). *Buying AI: Is the public sector equipped to procure technology in the public interest?* <https://www.adalovelaceinstitute.org/report/buying-ai-procurement/>
- Bode, I., & Huelss, H. (2025). Establishing human responsibility and accountability at early stages of the lifecycle for AI-based defence systems. *Ethics and Information Technology*.
- DAU. (2025). *Using AI to rapidly improve acquisition outcomes*. <https://www.dau.edu/4edacm/newsletters/jan-2025/using-artificial-intelligence>
- DOW. (2026). *AI strategy for the Department of War*. <https://media.defense.gov/2026/Jan/12/2003855671/-1/-1/0/ARTIFICIAL-INTELLIGENCE-STRATEGY-FOR-THE-DEPARTMENT-OF-WAR.PDF>



- European Parliament and Council. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. <https://data.europa.eu/eli/reg/2024/1689/oj>
- FNN. (2025). *AI to the rescue of federal procurement?* <https://federalnewsnetwork.com/federal-insights/2025/09/ai-to-the-rescue-of-federal-procurement>
- GAO. (2024). *Artificial intelligence: AI accountability framework and related resources*. <https://www.gao.gov/artificial-intelligence>
- GAO. (2025). *Generative AI use and management at federal agencies*. <https://www.gao.gov/products/gao-25-107653>
- GAO. (2025). *Artificial intelligence: Federal efforts guided by requirements and advisory groups (GAO-25-107933)*. <https://www.gao.gov/assets/gao-25-107933.pdf>
- GAO. (2026). *Fraud and improper payments: Data quality and a skilled workforce are key to leveraging AI*. <https://www.gao.gov/products/gao-26-108850>
- GSA. (n.d.). *Buy AI — GSA OneGov AI acquisition resources and contracting vehicles*. <https://www.gsa.gov/technology/government-it-initiatives/artificial-intelligence/buy-ai>
- GSA. (2025). *Use of artificial intelligence at GSA (CIO Order No. CIO 2185.2)*. <https://www.gsa.gov/directives-library/use-of-artificial-intelligence-at-gsa>
- Hubbard, R. (2025). *How AI can accelerate defense acquisition and help U.S. forces keep pace with technology*. Federal News Network.
- iQuasar. (2026). *AI procurement in 2026: How federal agencies will use automation to evaluate and monitor contracts*. <https://iquasar.com/blog/ai-procurement-in-2026-how-federal-agencies-will-use-automation-to-evaluate-and-monitor-contracts>
- Johnson, J. (2026). *Can AI behave ethically during military crises? Preserving human moral agency*. *International Affairs*.
- Kolt, N. (2025). *Lessons from complex systems science for AI governance*. PMC - NIH. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12365527>
- MeriTalk. (2026). *GSA rethinks federal acquisitions as AI reshapes contracting*. <https://www.meritalk.com/articles/gsa-rethinks-federal-acquisitions-as-ai-reshapes-contracting>
- MITRE (2024). *Enhancing acquisition outcomes through leveraging of artificial intelligence*. <https://www.mitre.org/sites/default/files/2025-03/PR-24-0962-Leveraging-AI-Acquisition.pdf>
- NIGP. (2026, February 11). *When artificial intelligence buys for us — Transparency and accountability in AI-driven procurement*. <https://www.nigp.org/blog/when-ai-buys-for-usv>
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1)*. <https://doi.org/10.6028/NIST.AI.100-1>
- National Institute of Standards and Technology. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile (NIST AI 600-1)*. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- Office of Management and Budget. (2024, March 28). *Advancing governance, innovation, and risk management for agency use of artificial intelligence (M-24-10)*. Executive Office of the President. <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf>



- Office of the Under Secretary of Defense for Acquisition and Sustainment. (2020, January 23). *DoD Instruction 5000.02 Operation of the Adaptive Acquisition Framework*. Department of Defense.  
<https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500002p.PDF>
- Open Contracting Partnership. (2025, November 10). *Governments risk “expensive failures” without smarter AI procurement, new report warns*. <https://www.open-contracting.org/news/governments-risk-expensive-failures-without-smarter-ai-procurement-new-report>
- Shipley Associates. (2026). *The human touch returns: How thoughtful AI integration is reshaping federal procurement*. <https://www.shipleywins.com/blogs/the-human-touch-returns-how-thoughtful-ai-integration-is-reshaping-federal-procurement>
- The White House. (2025). *White House unveils America’s AI Action Plan*. <https://www.whitehouse.gov/articles/2025/07/white-house-unveils-americas-ai-action-plan/>
- Tillipman, J. (2026). Abdicated judgment: AI tools and the future of reasoned decision-making in federal procurement. *Yale Journal on Regulation Notice & Comment*.  
<https://www.yalejreg.com/nc/abdicated-judgment-ai-tools-and-the-future-of-reasoned-decision-making-in-federal-procurement-by-jessica-tillipman>
- U.S. DoS. (2024). *Department of State compliance plan for OMB Memorandum M-24-10*.  
<https://2021-2025.state.gov/wp-content/uploads/2024/09/DOS-Compliance-Plan-with-OMB-M-24-10-Accessible-9.23.2024.pdf>





# AI Collusion in Procurement Auctions

**Andrew Tai<sup>1</sup>** – is an Assistant Professor of Economics at the Defense Security Cooperation, Department of Resources Management, located in Monterey, CA. He obtained a PhD in Economics from UC Berkeley and a BA in Economics and Applied Mathematics, also from UC Berkeley. [andrew.a.tai.civ@mail.mil and andrew.tai@nps.edu]

## Abstract

Auctions are an important format for defense procurement, historically accounting for at least \$1 billion per year (DoD Inspector General, 2012; GAO, 2018). A common format is the reverse auction, in which potential sellers submit bids for a contract, and the buyer selects the lowest bid (Alper & Boning, 2003; Coughlan et al., 2008). This design is well-suited for small uniform contracts where price is the only dimension of choice; common applications include commodities and routine services.

However, this setting is also conducive to the application of AI bidding algorithms by sellers. Such algorithms generally seek to maximize profit for their owners. An obvious way to do this is to collude. AI algorithms for bidding introduces the possibility that competing algorithms collude to increase prices without explicitly communicating with each other.

I conduct simulations in a simplified environment to test this possibility. I find that Q-learning algorithms indeed converge to supra-competitive prices in reverse auctions, allowing firms to extract higher than competitive profit. The learned policies also suggest they sustain collusion by punishing competitors for deviations.

## Introduction

Increasingly, firms use artificial intelligence (AI) algorithms to price goods and services. As early as 2015, Chen et al. (2016) find that one-third of sellers on Amazon had automated their pricing. Aggarwal et al. (2024) document extensive use of algorithms to automate bidding for advertising space in online auctions.

Auctions are an important format for defense procurement, historically accounting for at least \$1 billion per year (DoD Inspector General, 2012; GAO, 2018). A common format is the reverse auction, in which potential sellers submit bids for a contract, and the buyer selects the lowest bid (Alper & Boning, 2003; Coughlan et al., 2008). This design is well-suited for small uniform contracts where price is the only dimension of choice; common applications include commodities and routine services.

However, this setting is also conducive to the application of AI bidding algorithms by sellers. Indeed, the similarities between such applications and online advertising space are clear. While details are private, bidding algorithms generally seek to maximize profit for their owners. An obvious way to do this is to collude. The application of AI algorithms for bidding introduces the possibility that competing algorithms learn to collude to increase prices, even without explicitly communicating with each other.

In general, under U.S. antitrust law, collusion is defined as agreements between competitors to fix prices or market shares, requiring some “meeting of the minds” (U.S. Department of Justice, 2023). This paper takes a broader definition of “collusion” as including tacit behavior to fix prices at supra-competitive levels. Since the AI algorithms do not explicitly

---

<sup>1</sup> Contact: [andrew.a.tai.civ@mail.mil](mailto:andrew.a.tai.civ@mail.mil) and [andrew.tai@nps.edu](mailto:andrew.tai@nps.edu). Halligan Hall, M9. Disclaimer: The views expressed are those of the author and do not reflect the official policy or position of the Department of War or the U.S. Government.



coordinate with each other, it may be impossible to establish that the law was violated. Yet the outcome is the same—higher profits and thus higher costs to the buyer.

A common baseline learning algorithm is Q-learning. Prior research, including Calvano et al. (2020) and Klein (2021), has shown Q-learning algorithms learn to collude to extract profit above competitive prices. In the benchmark model of competition, Bertrand competition, the general pattern is Edgeworth cycling, characterized by an initial jump to a high price then a slow decline to the competitive price. During the decline period, competing firms receive higher profits due to the above-competitive price.

The possibility of autonomous collusion in procurement auctions was an open question. The literature has not been extended to reverse auctions. Nor had it been extended to more advanced learning algorithms such deep Q-learning, which is likely closer to state-of-the-art algorithms and can better accommodate complicated environments with fine price grids. This paper applies Q-learning to simulations of reverse auctions to evaluate whether and to what extent such algorithms learn to collude. I find that Q-learning algorithms indeed converge to supra-competitive prices in reverse auctions, allowing firms to extract higher than competitive profit. Depending on the experimental parameters, the profit extracted from the seller can be as much as 3x the competitive profit. The learned policies also suggest they sustain collusion by punishing competitors for deviations.

The results have implications for the cost and viability of auctions for procurement. If AI algorithms can collude to a high extent, the government may face increasing costs for contracts procured through auctions and similar formats. A theoretical understanding of these processes can support practical policy solutions to evolving challenges in defense procurement. Further results demonstrate the value of expanding the defense industrial base. Increasing the number of firms (even if all of them use AI algorithms for bidding) can markedly decrease the profit extracted. This paper's results also contain lessons beyond of the use of AI algorithms. Humans<sup>2</sup> can also replicate the behavior arising in these simulations, illustrating how tacit collusion can arise without explicit communication.

I also note this paper's limitations. First, this paper is an experiment—the findings here establish theoretical possibility. In reality, firms (both using AI algorithms and not) may or may not collude. Detecting such collusion may be difficult; I refer readers to Asker (2010) for details. Second, even though Q-learning agents converge to collusive pricing, doing so can take many periods, consistent with findings in Calvano et al. (2020). In real settings, these learning periods may result in large losses. However, more advanced algorithms may converge faster. Finally, the environment is highly stylized. The model is a major simplification of any real auction environment. However, the simplicity of the model presented and simulations offers clear and easily interpretable evidence of the possibility of algorithmic collusion.

The paper proceeds as follows. The Q-learning section recounts basics on Q-learning. The Experiments section documents the simulated experiments. The Results section gives the empirical results, and the Conclusion concludes.

## Q-learning

This paper trains Q-learning models to auction environments. While algorithms actually used by firms are likely much more advanced, little is publicly known. However, Q-learning is a popular basic algorithm. Advantages are that it can be defined using few parameters and that its behavior is readily interpretable. To my knowledge, Waltman and Kaymak (2008) are the first to

---

<sup>2</sup> Humans can of course be seen as very advanced general purpose learning algorithms.



apply Q-learning to an economic competition model. This section recounts basics of Q-learning; familiar readers may safely skip to the next section.

Q-learning is a model-free reinforcement learning (RL) algorithm. In a repeated environment, the Q-learning algorithm records payoffs from past actions and states to learn a state-action correspondence. This correspondence is typically referred to as a policy. Note that this does not refer to policy in the sense of principles or laws of a government. Through the rest of this paper, except in the conclusion where it is clear, policy refers only to the technical terminology for Q-learning algorithm behavior.

The original setup, due to Watkins (1989), is a stationary Markov decision process (MDP). Notably, Q-learning is for single agent decision environments. In each period  $t = 0, 1, 2, \dots$ , the agent observes a state  $s_t \in S$  and chooses an action  $a_t \in A(s_t)$ , where the available actions perhaps depend on the state. The agent receives reward  $\pi_t(s_t, a_t)$ . The next period's state is drawn from a probability distribution  $f(s_{t+1} | s_t, a_t)$ , perhaps dependent on the current action and state. The agent seeks to maximize expected net present value

$$\mathbb{E} \left[ \sum_{t=0}^{\infty} \delta^t \pi_t(s_t, a_t) \right]$$

where  $\delta < 1$  is the discount factor.

Maximizing this quantity is a dynamic programming problem, which is often solved analytically or computationally using the Bellman equation.

$$V(s_t) := \max_{a_t \in A(s_t)} \{ \mathbb{E}[\pi_t | s_t, a_t] + \delta \mathbb{E}[V(s_{t+1}) | s_t, a_t] \}$$

Intuitively, the value of a state is the combination of the reward from the action this period and the discounted value of the subsequent state, given the agent takes the optimal actions.

The  $Q$  function is similar but lists values of pairs  $(s, a)$ .

$$Q(s_t, a_t) = \mathbb{E}[\pi_t | s_t, a_t] + \delta \mathbb{E} \left[ \max_{a_{t+1} \in A(s_{t+1})} Q(s_{t+1}, a_{t+1}) | s_t, a_t \right]$$

Since the MDP is stationary, the  $t$  subscripts can be generic:

$$Q(s, a) = \mathbb{E}[\pi | s, a] + \delta \mathbb{E} \left[ \max_{a' \in A(s')} Q(s', a') | s, a \right]$$

where  $s$  is the current state, and  $s'$  is the next state.

When  $S$  and  $A$  are finite,  $Q$  can be represented as a matrix, where the  $Q_{s,a}$  entry is the expected reward for  $(s, a)$ . If the  $Q$  matrix is known, selecting the optimal policy is straightforward. Of course, the  $Q$  matrix may be unknown and difficult to solve for analytically. The Q-learning algorithm estimates the  $Q$  matrix over periods of repetition.

For  $t = 0$ , the  $Q$  matrix is initialized at  $Q_0$  arbitrarily (e.g. all 0s). Given  $Q_t$  and  $(s, a)$ , the algorithm updates

$$Q_{t+1}(s, a) = (1 - \alpha) Q_t(s, a) + \alpha \left[ \pi_t + \delta \max_{a' \in A} Q_t(s', a') \right]$$

and the non- $(s, a)$  entries remain the same. The parameter  $\alpha \in (0, 1)$  is the learning rate. Intuitively, the  $Q$  matrix is updated  $\alpha$  towards the current reward and discounted future value.

It remains to actually select  $a$  given  $s$ . This is the exploration policy, which is exogenously specified. Generally, Q-learning algorithms select the maximizing action with some probability, and randomly with the remaining probability. The simplest exploration policy is  $\epsilon$ -greedy exploration: with  $1 - \epsilon$ , the maximizing action (according to the current  $Q$  matrix) is chosen; with  $\epsilon$  probability, another action is drawn uniform randomly. More sophisticated policies incorporate exploration decay, where the chance of exploration decays to 0 as  $t \rightarrow \infty$ .



Under exploration decay policies, Q learning is guaranteed to converge to the optimal policy, shown by Watkins and Dayan (1992).

The tradeoff inherent in Q-learning is that exploration can be costly, yielding suboptimal rewards, but is necessary to learn the  $Q$  matrix to maximize future rewards.

Q-learning was originally designed for MDPs, which are single-agent and stationary environments; dynamic games are neither. In dynamic games, payoffs and state transitions may depend on both players' actions. Further, strategies may depend on the whole stream of past actions.<sup>3</sup> In full generality a state is thus never repeated; thus it is impossible to learn the infinite  $Q$  matrix. In this paper, the Q-learning algorithms treat the  $t - 1$  period actions as the state variable, and the game structure will not depend on the past. Finally, Q-learning is not guaranteed to converge to optimal actions as in MDPs. However, this leaves open the possibility of interesting behavior (and interesting research questions).

## Experiments

### Model

The economic environment considered is a reverse auction with simultaneous bids and two players. Each stage is one such auction, which is repeated  $T$  times.

Let  $i = 1, 2$  be the competing firms. In each period  $t \in [1, T]$ , the buyer (the “government”) offers a contract via a reverse auction with simultaneous bids. Each firm submits a bid  $b_{it} \in \{0, \frac{1}{k}, \dots, \frac{k-1}{k}, 1\}$ , where  $k$  is a step-size parameter. The profits to the firms are given by

$$\pi_{it}(b_{it}, b_{jt}) = \begin{cases} b_{it} - c & \text{if } b_{it} < b_{jt} \\ \frac{1}{2}(b_{it} - c) & \text{if } b_{it} = b_{jt} \\ 0 & \text{if } b_{it} > b_{jt} \end{cases}$$

That is, the lower bidder wins the contract, gaining his bid  $b_{it}$  minus the cost of providing the service  $c$ . If the two firms tie, the contract is split. (Equivalently, can represent randomization of the winner with expected utility firms.) In this paper, I normalize  $c = 0$ , though stochastic costs may be an interesting future research direction.<sup>4</sup> Notice that  $b_{1t} = b_{2t} = \frac{1}{k}$  is a Nash equilibrium profile. The firms collectively extract  $\frac{1}{k}$  profit per period.

I set a ceiling on bids at  $b_{it} = 1$ . The reasoning for this is practical rather than theoretical—the algorithms we consider require a finite action space. However, we can motivate this feature of the model either as a “reservation bid,” where the government posts that it will cancel the contract if the price is above 1, or by the presence of a third firm that will always bid on the contract at a price of 1.

### Q-learning Setup

To specify the Q-learning algorithm, we need the state variables and exploration policy.

To keep the model simple, I set the state to be the previous period's bids:

$$s_t = (b_{1,t-1}, b_{2,t-1})$$

<sup>3</sup> For example, consider subgame perfect Nash equilibria grim trigger strategies. The strategy punishes the opponent if he has ever defected, not just in the prior period.

<sup>4</sup> See Ballesterio (2025) for an example.



This allows the Q-learning algorithms to condition the present period's action  $a_t$  on the prior period's bids. Then  $|A| = k + 1$  and  $|S| = (k + 1)^2$ . Contrasted to a more complicated state variable, this restricts the complexity of the algorithms' policies. The primary advantage is faster learning, as there are fewer states to explore.

The exploration policy is  $\varepsilon$ -greedy with decay defined by

$$\varepsilon_t = \beta^t$$

where  $\beta > 0$  is a decay parameter. In  $t = 0$ , the algorithm chooses randomly. As  $t$  increases, exploration becomes less likely.

## Experimental Setup

I first summarize the parameters and note the baseline values for the simulations. To summarize, the parameters are:

- $T = 100,000$ , the number of repeated stages
- $k = 6$ , the discrete price step parameter
- $\alpha = 0.1$ , the learning rate
- $\beta = 0.99995$ , the exploration decay
- $\delta = 0.95$ , the discount factor

$T$  is chosen to be computationally feasible. The step size  $k$  is chosen to match the setting of Maskin and Tirole (1988). The learning parameter  $\alpha$  matches common benchmarks in computer science literature (c.f. Calvano et al., 2020). The exploration decay  $\beta$  is chosen to yield  $\sim 1\%$  chance of exploration after at  $t = 100,000$ ; that is,

$$\beta^{100,000} = 0.01$$

The tradeoff in selecting  $\beta$  is the necessity of visiting each  $Q$  matrix cell enough times to learn an optimal policy, versus losses resulting from suboptimal actions. Heuristically, with decay  $\beta$  and large  $T$ , there are in expectation  $\approx \frac{1}{1-\beta}$  exploration actions. At our standard parameters, there are  $(k + 1)^2 = 49$  cells in the  $Q$  matrix and 20,000 exploration actions; in expectation each cell is visited 408 times via random exploration.

Each iteration (of  $T = 100,000$  periods) constitutes one experimental observation. I repeat this setup 100 times and report results.

## Results

I first show bids and profits from the first experimental observation.

Players' bids increase to around  $b_{it} = 2/3$ , significantly above the competitive bids of  $b_{it} = 1/6$ . Correspondingly, the firms collect  $\pi_{it} = 1/3$  each, a combined profit of  $2/3$  out of the maximum profit 1. Figure 1 shows bids over time; Figure 2 shows profits over time. The two agents largely mirror each other, starting around the middle of the bid grid. Interestingly, one agent appears to dip towards Nash equilibrium before increasing toward  $2/3$ . I interpret this behavior as tacit collusion, as both agents sustain bids roughly 4x competitive bids.








**Figure 1. Bids Over Time**




**Figure 2. Profits Over Time**

Figure 3 displays the final learned policies. These are also suggestive of collusive behavior. For example, notice Agent 1's row for  $b_{1,t-1} = 2/3$ . If Agent 2 undercut this bid at  $b_{2,t-1} = 1/2$ , Agent 1 retaliates by bidding  $b_{1t} = 1/6$ . Though this evidence is only suggestive, it resembles limited punishment strategies in repeated games.




**Figure 3. Learned Policies**

### Average Results

To demonstrate that the pattern presented is not coincidental, I repeat the experiment 100 times. Figure 4 shows the average per-period bids across 100 repetitions. Unsurprisingly, the two agents' averages are almost identical. By  $t = 100,000$ , the Q-learning algorithms on average bid around  $b_{it} = 0.6$ , well above the competitive bids of  $b_{it} = 1/6$ . Figure 5 shows the 90% empirical confidence interval for a single agent's bids; that is, among the 100 repeated experiments, the 5th to 95th percentile per-period bids. These results demonstrate that supra-competitive behavior is remarkably consistent. Even at the 5th percentile, Firm 1 is bidding  $b_{1t} = 0.5$  by the final period.




**Figure 4. Bids Over Time, Average of 100 Simulations**




**Figure 5. Bids Over Time, Average of 100 Simulations, 90% Empirical Confidence Interval**

### Increasing the Number of Firms

I also introduce a third firm. The environment remains the same and completely symmetric. That is, the lowest bidder wins, and firms split profits uniformly among winners if there is a tie. As Figure 6 shows, supra-competitive pricing is preserved. However, the magnitude is reduced. By the last period on average, firms bid  $b_{it} = \frac{2}{6}$ , still above the competitive equilibrium of  $b_{it} = \frac{1}{6}$ , though greatly reduced from 0.6.




**Figure 6. 3 Firms, Bids Over Time, Average of 100 Simulations, 90% Empirical Confidence Interval**

This result highlights the importance of the number of firms. While pricing remains supra-competitive, it appears that increasing the number of firms decreases the profit extracted by the sellers.

### Conclusion

In this paper, I simulate a simplified sealed bid reverse auction environment, which governments often use to procure goods and services. I find that Q-learning algorithms consistently learn to collude, extracting multiple times the competitive equilibrium profit. On average, the firms extract three times the competitive profit. While these results are theoretical



rather than empirical, they have major implications for the use of auctions and similar formats to procure goods and services. In this paper's simplified environment, the government pays on average three times the competitive price by the last period.

Since the algorithms do not explicitly communicate with each other, such behavior would not generally be considered antitrust violations. Indeed, Q-learning algorithms simply learn to maximize profit; they do not receive any explicit instructions to employ anticompetitive behavior. I stop short of recommending particular government policy responses, though this paper's results demonstrate the importance of developing policies to address AI collusion. Such policies can include both explicit changes to the law to restrict use of algorithms and changes to the auction environment to reduce their collusive behavior. In this paper, I have simulated the auction environment with particular parameters—it may be helpful future work to examine behavior across parameters, especially the number of firms; such results may help determine what features of the environment promote algorithmic collusion.

## References

- Aggarwal, G., Badanidiyuru, A., Balseiro, S. R., Bhawalkar, K., Deng, Y., Feng, Z.,... & Zuo, S. (2024). Auto-bidding and auctions in online advertising: A survey. *ACM SIGecom Exchanges*, 22(1), 159–183.
- Alper, O. & Boning, W. B. (2003). *Using procurement auctions in the Department of Defense*. CNA Corporation. CRM D0007515.A3.
- Asker, J. (2010). A study of the internal organization of a bidding cartel. *American Economic Review*, 100(3), 724–762.
- Ballesterio, G. (2025). *Algorithmic collusion under sequential pricing and stochastic costs*. Available at SSRN 5203084.
- Calvano, E., Calzolari, G., Denicolo, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–3297.
- Chen, L., Mislove, A., & Wilson, C. (2016, April). An empirical analysis of algorithmic pricing on amazon marketplace. In *Proceedings of the 25th international conference on World Wide Web* (pp. 1339–1349).
- Coughlan, P., Gates, W., & Lamping, J. (2008). *Innovations in defense acquisition auctions: Lessons learned and alternative mechanism designs*. Acquisition Research Program.
- GAO. (2018). *Report GAO-18-446*.
- Klein, T. (2021). Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, 52(3), 538–558.
- Maskin, E., & Tirole, J. (1988). A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica: Journal of the Econometric Society*, 571–599.
- U.S. Department of Justice. (2023.) *Federal antitrust crime: A primer for law enforcement personnel*. <https://www.justice.gov/atr/page/file/1091651/dl?inline>.
- U.S. DoD Inspector General. (2012). *DODIG-2012-098*.
- Waltman, L., & Kaymak, U. (2008). Q-learning agents in a Cournot oligopoly model. *Journal of Economic Dynamics and Control*, 32(10), 3275–3293.
- Watkins, C. J. C. H. 1989. *Learning from delayed rewards* [Doctoral dissertation, King's College].
- Watkins, C. J., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3), 279–292.



# **An Actionable Framework & Interactive Web Application for Transitioning Technology to the Warfighter: The Transition Maturity Framework (TMaF) and Tech2PEO/PAE**

**Debra Zides**—is a principal program manager and agile acquisition expert at MITRE. A retired U.S. Air Force (USAF) lieutenant colonel acquisition officer, she served as deputy director of the USAF Life Cycle Management Center's acquisition excellence directorate and is a previous DAU guest lecturer. Zides holds a BS in Mechanical Engineering from Rensselaer Polytechnic Institute and MS degrees in Organizational Management and Military Operations from George Washington University and Air University, respectively.

**Natalie Faucher**—is a principal economic and business analyst at MITRE. She is an expert in life cycle acquisition and cost analysis. Faucher has led enterprise and program-level analysis and planning, including cost estimation, risk analysis, analysis of alternatives, and business case analysis for DoD, NATO, and international and civil agencies. She holds a BS in Economics from Shepherd University and an MS in Information Technical Management from Brandeis University.

**Tricia Youngbull** —is a principal energy technologies engineer/scientist and manages MITRE's technology transition support for the Operational Energy-Innovation Directorate. Youngbull was issued eight patents in her various industry positions focused on product development and manufacturing of photovoltaics, thermoelectrics, and integrated circuit packaging. She holds a BS and PhD in Materials Science in Engineering from The Ohio State University and the University of Washington, respectively.

**Alex Cocco** —is a MITRE lead energy technologies engineer contributing to techno-economic and supply chain analysis of defense critical energy technologies. Before joining MITRE, he worked at the U.S. Army Research Laboratory. Cocco holds a BS and PhD in Mechanical Engineering from the University of Connecticut and an MBA from the University of Maryland.

**RuthAnne Darling** —is the Director of the Operational Energy-Innovation (OE-I) Directorate in the Office of the Under Secretary of War for Acquisition and Sustainment (OUSW(A&S)). She manages the Operational Energy Capabilities Improvement and Prototyping Funds, overseeing more than \$250 million in cutting-edge energy technology solutions. Darling holds a BS in Computer Science from the University of Massachusetts and an MS in Business Management from the University of Maryland.

## **Abstract**

MITRE developed the Transition Maturity Framework (TMaF) and Tech2PEO/PAE in partnership with the Operational Energy-Innovation Directorate (OE-I) within the Office of the Under Secretary of War for Acquisition and Sustainment (OUSW(A&S)). The TMaF offers program managers a standardized means to evaluate, balance, and monitor transition maturity. For innovators, it provides actionable guidance to inform strategic resourcing/execution of critical activities needed to navigate transition. Starting with its FY25 call for proposals, OE-I fully integrated the TMaF into its proposal solicitation, selection, and execution management processes. As OE-I's FY25 cohort executes their projects, periodically reporting progress via TMaF assessments, they are creating a longitudinal, technology maturation data set. This TMaF data stands to be a key resource to inform decision-makers on how best to invest in maturing technologies to increase the likelihood of their successful transition to warfighters. As OE-I's TMaF rollout continues, it is working with MITRE to create complementary transition tools, such as Tech2PEO/PAE. Tech2PEO/PAE is designed to help innovators work through a common barrier to transition: identifying and connecting with a relevant program. Both the TMaF and Tech2PEO/PAE are available via MITRE's Acquisition in the Digital Age site at <https://aida.mitre.org/tmaf>.

**Keywords:** technology transition; readiness; partnerships; acquisition; industrial base; assessment; risk; small business participation; acquiring emerging technologies; commercial integration





## Paper

The Department of War (DoW) invests billions of Research, Development, Test and Evaluation (RDT&E) dollars into emerging technologies to maintain U.S. military superiority over our adversaries. This funding supports technology development and transition out of the laboratory, into an acquisition program, and on to warfighter use. Too often, however, groundbreaking DoW technologies fail to advance to full-scale deployment because navigating transition requires DoW domain knowledge often outside the strengths of a typical Science and Technology (S&T) project team (e.g., manufacturing, cost estimation, defense acquisition, requirements, and operator training). This knowledge gap increases an S&T project's risk of falling in to the "Valley of Death" because they underestimate the time and resourcing necessary to execute critical activities to meet the technology acceptance criteria of essential transition stakeholders (e.g., acquisition program managers, the requirements community, and warfighters). Effective technology transition management therefore is critical to rapidly and cost-effectively delivering advanced capabilities to the warfighter.

Here we present two transition management tools created by a multidisciplinary team of warfighters and subject matter experts in defense acquisition and technology research and development: the Technology Maturity Framework (TMaF) and the Technology-to-Program Executive Office/Portfolio Acquisition Executive (Tech2PEO/PAE) application, both accessible at [aida.mitre.org/tmaf](https://aida.mitre.org/tmaf).

MITRE developed the TMaF and Tech2PEO/PAE in partnership with the Office of the Under Secretary of War for Acquisition and Sustainment's (OUSW(A&S)) Operational Energy-Innovation Directorate (OE-I; Jackson, 2025; Zides et al., 2025). OE-I invests in the development and transition of innovative operational energy solutions to meet joint requirements (OUSW(A&S) Operational Energy-Innovation Directorate, 2026). OE-I recognizes that the challenge of transitioning their technologies requires more than technical and manufacturing readiness. As a result, the TMaF is designed to provide an actionable toolkit to ensure a technology is not only technically proven and manufacturable, but also favorably positioned for acquisition, aligned to a funded requirement/operational problem, and progressing toward a state where warfighters are both willing and able to adopt it.

The TMaF (see Figure 17) consists of five dimensions, each measured on a 1–9 scale, that align with the technology acceptance criteria of key defense acquisition stakeholders: joint/service innovation execution organizations, industrial base production, defense acquisition program offices, the defense requirements/operational problems community, and warfighters. Of the five TMaF dimensions, Technology Readiness Level (TRL), Manufacturing Readiness Level (MRL), and Transition Confidence Level (TCL) were preexisting. TRL, MRL, and TCL primarily account for the priorities of innovation execution organizations, industrial base production, and acquisition programs, respectively, but do not fully capture the priorities of the requirements community and warfighters. To close this gap, MITRE and OE-I developed the Requirements and Warfighter Readiness Levels (RRL and WRL). RRL assesses a technology's alignment to a validated/funded requirement or key operational problem. WRL measures the extent that warfighters are willing and able to adopt a technology. It evaluates solution adoption in a manner aligned with the fundamental tenants of Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, Facilities, and Policy (DOTMLPF-P) analysis. Together, the TMaF dimensions provide a standardized means to think "beyond TRL" and manage transition holistically through thoughtful alignment with requirements/key operational problems and early, meaningful engagements with warfighters and acquisition programs.






**Figure 17. The Transition Maturity Framework**

Throughout its development, OE-I emphasized the importance of the TMaF being actionable. Thus, for each TMaF readiness dimension and level, the framework provides representative critical activities needed to progress to the next level. These activities offer innovators a starting point to develop a strategic roadmap for navigating their project through the transition milestones captured by the TMaF levels. For example, to progress from WRL 5 to WRL 6, the TMaF provides four critical activities, including “involve warfighters in shaping policies and procedures to ensure practicality and acceptance.” This example exhibits the three essential elements of TMAF critical activities. It indicates which stakeholder to engage (“warfighters”), when to engage them (“when shaping policies and procedures”), and the engagement’s goal (“ensure practicality and acceptance”). Through its combination of levels and critical activities, the TMaF provides a standardized means to baseline, act upon, track, and certify a technology’s maturation that enables innovators and RDT&E portfolio managers to understand a technology’s maturity level and strategically allocate resources to manage its transition. Tables detailing TMaF levels and critical activities are available via MITRE’s Acquisition in the Digital Age website’s TMaF page (<https://aida.mitre.org/tmaf>).

Starting with its FY25 call for proposals, OE-I established TMaF entry/exit levels for their programs and fully integrated the TMaF into its proposal solicitation, selection, and execution management processes. Funded project teams use the TMaF to communicate transition progress to OE-I and to advance their technology’s transition through execution of critical activities such as targeted stakeholder outreach, partnerships (e.g., with the industrial base), and warfighter engagements/demonstrations. OE-I uses the TMaF data generated by project teams for Congressional reporting in the Operational Energy Management and Innovation (OMNI) tool, an NDAA-mandated database for tracking transition of OE technologies across the DoW.



As OE-I continues to integrate the TMaF into its organizational processes, it is proactively working with MITRE to collect/address user feedback and to develop complementary tools for reducing transition barriers, such as Tech2PEO/PAE. Tech2PEO/PAE is an interactive web application that helps technology developers connect with acquisition organizations. Users input a short description of the technology and, if desired, their technology's corresponding TMaF assessment. The software then generates a curated list of PEOs, and soon, Portfolio Acquisition Executives (PAEs), their associated contact information, and relevant funding opportunities. Interested users can register for a Tech2PEO/PAE account by emailing [tech2peo@mitre.org](mailto:tech2peo@mitre.org). The TMaF stands apart from earlier technology readiness frameworks with holistic, actionable guidance for transition management, incorporating two new readiness dimensions, and integrating complementary tools such as Tech2PEO/PAE to reduce transition barriers.

As the first OE-I cohort using TMaF executes their projects, periodically reporting progress via TMaF assessments, they are creating a longitudinal data set. Already, OE-I is establishing the analytic infrastructure necessary to capture and fully leverage this data. This kind of data could be a key resource to RDT&E portfolio managers and the broader defense acquisition ecosystem. It offers the opportunity for analysis of, for example, a program manager's return on investment (based on the resourcing required to execute TMaF critical activities); transition management at portfolio and project levels; and key determinants of transition success, which can inform/motivate the development of complementary tools like Tech2PEO/PAE to address transition barriers. These TMaF-enabled analytics have the potential to inform decision-makers on how best to invest in maturing emerging technologies to increase the likelihood of successful transition to warfighters.

Get started with the TMaF and Tech2PEO/PAE today via the resources linked in Figure 18 and connect with questions/feedback by emailing [tmaf@mitre.org](mailto:tmaf@mitre.org)!



**Access** TMaF and Tech2PEO/PAE through MITRE's Acquisition in the Digital Age (AiDA) website

**Listen** to a Federal News Network Interview on TMaF with MITRE MNS SVP & GM, Keoki Jackson

**Read** our Defense Acquisition Magazine Article

Figure 18. Engage with the TMaF and Tech2PEO/PAE via Resources Hosted at <https://aida.mitre.org/tmaf>

## References

- Jackson, K. (2025, October 28). *A new framework aims to build a bridge across the defense acquisition valley of death*. T. G. Federal News Network, Interviewer.
- OUSW(A&S) Operational Energy-Innovation Directorate. (2026). *Operational energy management system*. <https://oecif.org/>
- Zides, D., Faucher, N., Youngbull, T., Cocco, A., & Darling, R. (2025, September–October). A practical framework for transitioning technology to the warfighter. *Defense Acquisition Magazine*, 14–18.



## PANEL 7. THE ARSENAL OF INNOVATION: REVITALIZING INDUSTRIAL SURGE AND MANUFACTURING SECURITY

Wednesday, May 6, 2026, 1315 – 1430 ET (1015 - 1130 PT)

### Panel Summary:

Global deterrence and warfighting success depend on an industrial base capable of rapid mobilization and sustained production. This panel addresses the erosion of U.S. industrial might and explores frameworks to rebuild surge capacity through strategic government-industry collaboration. By examining the creation of a Manufacturing Security Program (ManuSP), analyzing high-stress case studies like the Civil Reserve Air Fleet, and investigating chronic underinvestment in critical sub-tier industries like aerospace gears, these researchers offer a roadmap to transition from fragile, just-in-time supply chains to a robust "Arsenal of Democracy" suited for modern peer competition.

**Chair: Todd Lyons, Col USMC (Ret.),** Vice President for the Naval Postgraduate School Alumni Association and Foundation

### Presenters:

**Connecting the Kill Chain to the Supply Chain: Building Industrial Surge Capacity – Dr. John Matlik, Fellow at Northrop Grumman Aeronautics Systems**

**Establishing a Manufacturing Security Program (ManuSP) – Katy Buda, Center for Strategic and International Studies (CSIS)**

**The Gears of War: The State of the Aerospace Gear Industry and the Implications for Accelerating Acquisition – LtCol Nathaniel L. Ross, USMC, PEO Air Anti-Submarine Warfare, Assault & Special Mission (PEO(A))**



**Todd Lyons, Col USMC (Ret.)**—Retired Colonel Todd Lyons serves as the Vice President for the Naval Postgraduate School Alumni Association and Foundation. In this role, he bridges the divide between industry, academia, and DoD entities to accelerate the responsiveness of NPS to the challenges arising from great power competition and emerging technology. Lyons also serves as a volunteer instructor for Innovation Leadership at NPS.

Lyons served in the Marine Corps for 30 years. In his last assignment, he served as the Senior Marine Representative and the Associate Dean of Research at the Naval Postgraduate School. In those roles, he connected the work of students and faculty to the operational challenges of the Marine Corps, the Navy, and the broader DoD. As an innovation leader, he leveraged the research at NPS and the operational experiences of the students to drive the adoption of new practices across the Naval Services.

After graduating from the Naval Postgraduate School in 2000, Lyons became a Middle East Foreign Area Officer and an Intelligence Officer with assignments at the Marine Corps Intelligence Activity and the Defense Intelligence Agency. Todd spent nine of his last eighteen years in the Marine Corps serving in a variety of assignments across the Middle East, including Qatar, Israel, Oman, and Iraq.

Lyons earned a Master of Arts with honors in National Security Affairs from the Naval Postgraduate School. He earned a Master of Science degree from the University of Haifa and was the Honor Graduate from the Israeli National Defense College in 2007. For his reporting during the Israeli withdrawal from Gaza in 2005, he received the Department of State Superior Honor Award. He received a second Superior Honor Award from the Department of State for his service in Oman from 2011-2014.

Lyons is married and has three daughters. He enjoys being active and keeping up with his family.



# Connecting the Kill Chain to the Supply Chain: Building Industrial Surge Capacity

**Michael McGrath, DSc**—is a consultant for the University of Maryland Applied Research Lab for Intelligence and Security (ARLIS), a Research Fellow at the Acquisition Innovation Research Center (AIRC), and a member of the NDIA Ad Hoc Industrial Surge Working Group. [mmcgrat1@stevens.edu]

**John Matlik, PhD**—is a Fellow at Northrop Grumman Aeronautics Systems, an American Institute of Aeronautics and Astronautics (AIAA) Associate Fellow, lifetime National Defense Industry Association (NDIA) member, and chair of the NDIA Ad Hoc Industrial Surge Working Group. [John.Matlik@ngc.com]

**Olivia Pinon Fischer, PhD**— is a Principal Research Engineer at the Georgia Institute of Technology, an American Institute of Aeronautics and Astronautics (AIAA) Associate Fellow, and the chair of the AIAA Digital Engineering Integration Committee. [olivia.pinon@asdl.gatech.edu]

## Abstract

U.S. industrial might, once a decisive deterrent in the post–World War II era, has significantly eroded. Consolidation of the defense industrial base, offshoring of commercial manufacturing, and brittle just-in time supply chains have hollowed out our ability to deter or defeat peer adversaries. This research builds on scenario planning performed by the National Defense Industrial Association (NDIA) Surge Capacity working group. It develops a framework for actions and investments during peacetime for an industrial mobilization capability—one that is prudent to plan and hopefully never needed. The research examines quantification of demands for ordnance, spares and maintenance at combat utilization rates; resource strategies for building the needed industrial capacity for a prolonged conflict; acquisition strategies for pre-planned surge capacity; sustainment actions to maintain supply chain visibility; transition strategies that re-create an industrial mobilization board; and rapid response adaptation strategies for combat contingencies. A targeted literature review evaluates prior planning efforts, organizations and models dating back to the Cold War. The framework proposes courses of action suited to today's defense and commercial industrial base (including allies).

## Introduction

The Department of War (DoW) Acquisition Transformation Strategy seeks to “Put the entire acquisition system and the industrial base on a wartime footing” (DoW, 2025a). As this strategy is implemented, the Department will be increasingly prepared to respond to the demands of a short, intense conflict with a peer adversary. Efforts are underway to rapidly rebuild the inventory of munitions needed for such a conflict. But more needs to be done to prepare for the possibility of a prolonged war. Today's Defense budget is around 3.5% of Gross Domestic Product (GDP) compared to 40% of GDP during WWII (Figure 19). Ukraine in 2024 spent 34% of GDP on defense (Tian et al., 2024). We are far from being in a wartime economy and hope to remain so. It is prudent, however, to plan now for an industrial mobilization we hope never to need.





**Figure 19. U.S. Defense Spending as % of GDP**  
(Federal Reserve Bank of Saint Louis [FRED], 2026; O'Hanlon, 2024; Stockholm International Peace Research Institute, 2025)

This paper is based on the work of the National Defense Industrial Association (NDIA) Ad Hoc Working Group on Industrial Surge Capacity. That group examined unclassified scenarios in a contested logistics environment to demonstrate simulation-based methods to quantify replenishment demand profiles beyond the initial 60 days (NDIA Manufacturing Division, 2025). In the next section of this paper, industrial mobilization approaches dating from World War I are reviewed. The third section examines scenario planning and presents a methodology for quantifying demand profiles for different commodity classes. In the fourth section, the DoW Acquisition Transformation Strategy and Defense Industrial Base (DIB) ecosystem are reviewed as the key enablers for acquisition improvements that are vital for industrial surge and the foundation for pre-planning industrial mobilization. The final section develops a playbook with pre-planned actions (herein called “switches”) for realizing rapid industrial mobilization and offers recommendations for near-term action and longer-term research and development.

## **Background on U.S. Industrial Mobilization Planning**

The historical context for modern U.S. industrial mobilization starts with preparation for World War I. As described by McGinn (2020), in the early 1900s, the U.S. industrial base was focused entirely on commercial manufacturing, except for a limited number of arsenals and shipyards. Unlike European nations, where military planners mobilized industrial production of war materiel in anticipation of the coming conflict, the United States entered the war in 1917 largely unprepared. The War Industries Board (WIB) was created as an emergency agency staffed by industry executives on loan from their companies to oversee war production. The emergency mobilization built enough cargo ships, rifles, trucks and aircraft to support the deployment within 2 years. Most of this materiel arrived too late to equip the 2 million U.S. troops, who fought using heavy equipment from European allies. The WIB was dismantled in 1919, but the importance of U.S. industrial might and the lessons of mobilization planning had been learned.



During the 1930s, President Roosevelt saw the growing need for industrial preparedness. Congress passed the Fleet Expansion Act (Naval Act of 1938) to authorize a 20% increase in U.S. Navy strength and procurement of 3,000 aircraft, moving the Navy toward a two-ocean fleet. In 1940, President Roosevelt asked Congress initially for \$2 billion for defense, followed in May by a request for another \$2.4 billion, and then in July for an additional \$4.8 billion

To procure reserve stocks of tanks, guns, artillery, ammunition, etc., for ... a total of 2,000,000 men if a mobilization of such a force should become necessary. To provide for manufacturing facilities, public and private, necessary to produce critical items of equipment for a land force of 2,000,000 men, and to produce the ordnance items required for the aircraft program of the Army and Navy—guns, bombs, armor, bombsights and ammunition. Procurement of 15,000 additional planes for the Army and 4,000 for the Navy, complete with necessary spare engines, armaments, and the most modern equipment. (Roosevelt, 1940)

This was the basis for the “Great Arsenal of Democracy” discussed in Roosevelt’s December 1940 famous fireside chat.

The WWII whole-of-economy mobilization became the benchmark for wartime industrial planning. Leaders of industry, such as Bill Knudsen of General Motors and Henry Kaiser, a construction and shipbuilding magnate, were instrumental in establishing or converting to wartime production industrial capacity in steel, ships, tanks, aircraft, and all the materials needed for war. The Government brought in industry to help with the War Production Board, the Office of Production Management, the Supply Priorities and Allocation Board, and others. The Automotive Council for War Production (ACWP) and other industry groups pooled resources, tools and technical knowledge to increase efficiency, enabling mass production of B-24 bombers. To illustrate the magnitude of the transition to wartime production, there were about 3 million automobiles manufactured in the United States in 1941. During the entire war, only 139 additional cars rolled off the assembly lines. Instead, automakers built guns, trucks, tanks and aircraft engines. The Lionel toy train company started producing items for warships, including compasses. Alcoa, the aluminum company, produced airplanes. The Mattatuck Manufacturing Company, which had made upholstery nails, switched to making cartridge clips for Springfield rifles. Shipyards turned out entire fleets of aircraft carriers, battleships, destroyers, submarines and other vessels (Vergun, 2020). These efforts resulted in an unprecedented increase in production, with U.S. factories producing more than the combined output of all Axis powers by 1944.

In 1950, Congress enacted the Defense Production Act (DPA), which remains an important tool today for industrial preparedness. DPA Title I: Priorities & Allocations allows the government to “jump the line.” Agencies can require companies to prioritize government contracts over private ones. Title III: Expansion of Capacity provides financial incentives (loans, grants, or purchase commitments) to help private companies expand production. Title VII: General Provisions grants the government power to block foreign mergers that threaten national security and establishes “Voluntary Agreements” with industry. While Title VII is routinely used by the Committee on Foreign Investment in the United States (CFIUS), Voluntary Agreements have fallen into disuse and may become important in planning for future industrial mobilization (McGinn, 2025).

In recent years, the DPA authorities were used to enable the U.S. response to COVID-19, called Operation Warp Speed (OWS). The DPA allowed the U.S. government to accelerate vaccine development from the typical 10-year timeline down to less than 1 year. Rather than the government simply buying finished products, the DPA allowed officials to assert priority and manage the entire supply chain. As part of OWS, the Department of Defense (DOD) used Other



Transaction Authority (OTA) consortia as the primary legal and administrative vehicle to move billions of dollars to pharmaceutical companies at a speed that traditional federal contracting could not match. In essence, the DPA was the basis for expanding manufacturing and the OTA consortia were the fast-track lane for the money and legal agreements.

A comprehensive CSIS assessment identified several enduring insights from the history of industrial mobilization (Cancian et al., 2020). Among the most significant: Mobilization takes years and constantly finds itself behind the pace of demand; there is a persistent need for centralized economic planning to cope with the massive dislocations that arise; production must be balanced across supporting capabilities as well as major end items; and there is critical value in beginning industrial mobilization before conflict begins. There is also an inherent tension between efficient peacetime production and maintaining capacity for wartime surge, a tension that must be resolved through deliberate planning rather than ad hoc responses. These historical and contemporary realities provide the motivation for the mobilization framework developed in the following sections of this paper.

## Scenario Planning and Quantifying Demand

The NDIA Industrial Surge Working Group's Scenario Planning effort uses realistic, unclassified narratives (such as F-35 component failures, contested logistics, and affordable-mass drone deployments) to run iterative wargames and high-fidelity simulations with a wide range of defense, industry, and academic stakeholders, while calling for a parallel classified scenario set to capture detailed threat effects. Quantifying Demand focuses on the overwhelming wartime need for long-range munitions and accelerated spare part consumption, employing discrete event simulation and wargaming informed hybrid methods to produce time phased, mission coupled demand profiles. These profiles identify peak consumption periods and the specific "switches" that trigger surge contracts, tooling expansion, and workforce mobilization. Together, the two sections provide the analytical foundation for a preemptive "break glass" surge plan that aligns industrial capacity with dynamic logistics demands in high intensity conflict.

## Scenario Planning

Since its inception, the NDIA Industrial Surge Working Group has used realistic, fictional, and unclassified scenarios as a way to support the development of a surge framework and "kill-chain-to-supply-chain" playbook. By deliberately keeping the scenarios unclassified, the group was able to involve a broad spectrum of stakeholders, test innovative tools (such as the USAF Deployment Readiness Tool [Crowley, 2022]) and generate actionable recommendations.

The Working Group blended iterative wargaming and simulation, subject-matter-expert validation and cross-community engagement in a series of multi-day workshops. Workshop participants included representatives from the Office of the Under Secretary of Defense for Acquisition & Sustainment, DARPA, the Air Force Research Laboratory, the Manufacturing Innovation Institutes, leading defense OEMs, and lower-tier suppliers. Simulation iterations were performed by Georgia Tech to model demand signals, failure modes, and logistics bottlenecks, allowing participants to receive rapid feedback on how changes in the kill chain ripple through the logistical supply chain and industrial base.

Three scenarios were co-developed with Tim Zadalis (Major General, USAF, Ret.) and examined in this series of workshops:

- **F-35 Critical Components** – Considered battle damage to a structural wing spar, failure of an electronic stabilizer control, and a post mission recovery decision that diverts the aircraft to a maintenance hub based on real time health data. Highlighted how digital



twins might be used to generate demand spikes for specific high value parts, enabling the “switch on” of surge contracts for rapid response/retooling of supplier lines.

- Contested Logistics – Integrated model of Agile Combat Employment sorties supported by Air Mobility Command (AMC) airlift and Air Combat Command (ACC) kinetic missions, with a joint all domain command and control overlay that prioritizes combat sortie generation and how it influenced the supply chain and sustainment. Highlighted the interdependence of airlift and kinetic sortie rates, quantifying the “logistics-to-kill-chain” coupling that the surge playbook must address (e.g., surge capacity for rapid depot level repairs of subsystem and Line Replaceable Unit [LRU] components).
- Affordable Mass – Carrier-based deployment of attritable kinetic kill drones for area defense and Collaborative Combat Aircraft (CCA) formations of enhanced F 35/F 22 platforms executing high tempo sortie cycles. Provided context for evaluating low cost, high volume production pathways and the scalability of “affordable mass” surge lines that complement traditional high value platforms.

While these scenarios generated insights in an unclassified setting, classified scenarios that the DoW can share with cleared industry in an appropriate setting are now needed. These would allow inclusion of granular threat modeling (e.g., sensor, electronic-warfare, and cyber-attack effects) that directly impact failure rates and logistics node survivability; that otherwise cannot be considered and addressed in an unclassified environment, e.g., access to DoW-level demand forecasts, contract award data, and classified inventory positions, which are essential for calibrating “surge readiness level” metrics, and operational lessons learned from past DoW engagements and exercises. Future work should therefore establish a classified scenario-working group that mirrors the unclassified structure (with government lead, industry partners, and academic analysts) while operating under appropriate DoW security controls. The Space Force is already pursuing direct industry collaboration through a series of classified wargaming activities (Albon, 2026). The output will be a set of “protected” scenario packages that could leverage a classified simulation environment, allowing the unclassified surge framework and playbook to be stress-tested against realistic adversary actions and to refine the decision-point “switches” that trigger surge contracts, expanded workforce mobilization, and rapid tooling acquisition.

## Quantifying Demand

The most significant demand on the defense industrial base in a major war would likely come from munitions expenditures and the wear and tear of weapons systems and equipment (Jones, 2023). Analyses of high-intensity conflict scenarios consistently demonstrate that current stockpiles are insufficient (Gunziger, 2021; Jones, 2023). The CSIS wargame, which modeled a 2026 invasion scenario across 24 iterations using analytically based rules grounded in historical data and operations research, consistently surfaced the centrality of long-range precision munitions to the defense of Taiwan (Cancian et al., 2023; Cancian et al., 2020).

Lead times compound the problem. Missiles such as the PAC-2/PAC-3, Tomahawk Block V, JASSM, and Precision Strike Missile take about 2 years just to get first units off the line (Jones, 2023). The DoW Munitions Acceleration Council (DoW, 2026) aims to at least triple production capacity for critical munitions through multi-year contracts, but even with acceleration, restocking at wartime consumption rates will take years.

The challenge extends beyond munitions to spare parts and sustainment materiel. In high-tempo operations, aircraft and vehicles experience accelerated wear, driving demand for line replaceable units, structural components, and consumables at rates far exceeding



peacetime planning assumptions. Initial stocks of spares and consumables will need to be replenished during any extended combat, likely exceeding the capacity of peacetime supply chains.

Effective industrial surge planning requires more than aggregate estimates of materiel requirements. In high-intensity conflict against a peer adversary, demand is *time-phased* and *operationally coupled*: What matters is the shape of consumption over days and weeks and how that consumption propagates through maintenance pipelines, transportation networks, and inventory positions to influence the number of missions that can actually be executed. In other words, it is essential to recognize the operational dynamics and how materiel eventually translates into combat power/ability to fight. In essence, the supply chain is the kill chain.

The importance of this coupling between operational outcomes and industrial planning is amplified by the realities of contested logistics. Despite decades of analysis, logistics remains strategically decisive but rarely given equal footing with fires, maneuver, or command and control (Welch, 2026). Quantifying wartime demand requires methodological tools that can capture the complexity of operational consumption patterns and their interaction with logistics and supply chain constraints. Existing approaches span a range of fidelity and analytical focus:

- Static requirements estimation. The most basic approach involves calculating aggregate materiel requirements based on assumed operational tempos and historical consumption rates. While straightforward, this method treats demand as a fixed quantity and does not capture the dynamic interaction between consumption, resupply, and operational outcomes.
- Scenario-based operations research. More sophisticated approaches embed demand estimation within operational scenarios that model force employment, attrition, and logistics constraints. The CSIS wargame of a Chinese invasion of Taiwan exemplifies this approach, combining historical analogies (e.g., amphibious lift rates derived from Normandy, Okinawa, and the Falklands) with theoretical weapons performance data to adjudicate combat outcomes across 24 iterations (Cancian et al., 2023).
- Discrete event simulation. Discrete event simulation (DES) offers the highest fidelity for modeling time-phased demand because it can represent individual events—sorties flown, weapons expended, components failed, repairs completed, parts requisitioned—as they occur in sequence over simulated time. The DRT Tech Contested Logistics Operations Under Digital Support (CLOUDS) solution, employed in NDIA Working Group workshops, uses this approach to model demand signals, failure modes, and logistics bottlenecks in contested environments (Cantrell et al., 2025), allowing participants to observe how changes in the kill chain ripple through the logistical supply chain and into the industrial base (NDIA Manufacturing Division, 2025). Similarly, the U.S. Air Force-funded Deployment Readiness Tool (DRT) uses component reliability assessments to quantify spare parts demand under operational scenarios.
- Wargaming-informed hybrid approaches. A promising direction combines the strategic breadth of wargaming with the analytical precision of simulation. In this approach, wargame outcomes, e.g., force dispositions, attrition rates, operational tempos, serve as inputs to detailed simulation models that generate time-phased demand profiles for specific commodity classes. This hybrid methodology preserves the operational realism and expert judgment that wargames contribute while adding the quantitative rigor needed to translate operational outcomes into industrial planning signals.





The NDIA Industrial Surge Working Group developed a methodology that operationalizes the time-phased, mission-coupled demand concept through a structured, multi-step process. First, the operational scenario defines the mission profile: platform types, sortie rates, threat environment, and expected engagement patterns. Second, the scenario is instantiated in a discrete event simulation that models weapon expenditures, component failures, battle damage, and maintenance actions at the individual-event level over the full duration of the conflict. Third, the simulation generates time-phased demand profiles for multiple commodity classes, e.g., munitions (by type), spare parts (by system and line replaceable unit), consumables, and repair materials. These profiles capture not only the aggregate quantity demanded but the temporal distribution of that demand: peak consumption periods, sustained burn rates, and the interaction between munitions expenditure and maintenance-driven demand. Fourth, the demand profiles are compared against current production rates and inventory positions to identify the magnitude and timing of surge requirements. The gap between projected demand and baseline capacity defines the industrial response needed and the pre-planned “switches” that must be activated at the onset of conflict.

The United States cannot afford to wait until a conflict begins to determine how the defense industrial base should respond. Jones (2023) argued that the DoD, in coordination with Congress, should develop a plan that pre-positions the policy, contractual, and regulatory steps needed to streamline production, acquisition, replenishment, foreign military sales, and ITAR processes in an emergency wartime situation. The demand quantification methodology described above provides the analytical foundation for such a plan by identifying, in advance, the specific commodity classes, quantities, and timelines that drive the most critical surge requirements.

Finally, over 100 logistics-focused wargames and exercises since 2019 have consistently exposed the same contested-logistics vulnerabilities, i.e., findings that are by now well established. Yet these insights have not been translated into revised planning assumptions, force design changes, or binding resource decisions. The gap is not analytical but institutional: There is no mechanism to connect what wargames reveal to what programs fund and commanders execute. Closing that gap requires a decision-support ecosystem in which operationally relevant data, e.g., demand profiles, supply chain capacity, inventory positions, and throughput constraints, is continuously available, authoritative, and linked to the governance structures that allocate resources. Without that connective tissue, wargame findings will continue to accumulate without leading to positive outcomes.

## **A Playbook for Industrial Surge**

Translating these insights into an actionable playbook requires addressing three enabling conditions. First, contracting mechanisms must be in place before conflict begins. The DoW Acquisition Transformation Strategy paves the way for this. Second, workforce readiness is a binding constraint. Realistically, mobilization needs to tap the capacity of the U.S. commercial production base to meet production manpower needs. Third, supply chain visibility must extend deep into the defense and commercial (and allied) sub-tier base to identify capabilities and capacity that can avert wartime bottlenecks.

### **The DoW Acquisition Transformation Strategy Provides a Foundation**

The DoW Acquisition Transformation Strategy addresses long-standing shortcomings in the interface between the DoW and the industrial base. Its five “pillars” not only transform the acquisition process to a new emphasis on speed and warfighting capability but also provide a new ecosystem that can enable rapid transition to industrial mobilization if needed. The five pillars and the sub-elements most relevant for surge capacity planning are:



1. Fuel the Arsenal of Freedom: Rebuild the Defense Industrial Base
  - Expand the Industrial Base, Stabilize Demand Signals, Accelerate Private Capital Investment
  - Accelerate Commercial Preference, Procure Industry-Driven Solutions, Information Technology Acquisition
  - Establish the Industrial Base Consortium
2. Elevate and Empower the Acquisition Workforce to Rapidly Deliver Capability
  - Portfolio Acquisition Executives, Capability Portfolio Management, Capability Trade Councils
  - Accelerate Planning, Programming, Budgeting, and Execution (PPBE) Reform
3. Maximize Acquisition Flexibility through Reduced Regulations and Processes
  - Digitize Acquisition
  - Modernize Test Infrastructure, Reduce Test Oversight
  - Common-Sense Accounting, Federal Acquisition Regulation (FAR) and Defense Federal Acquisition Regulation Supplement (DFARS) Reduction, Foreign Military Sales (FMS) Reform
4. Develop High Performance Systems through Rigorous Enterprise Technical and Execution Excellence
  - Wartime Production Unit
  - Open Systems Architectures, Rapidly Generate New Kill Chains
5. Improve Effective Lifecycle Risk Management
  - Government Weapons Repair and Maintenance
  - Fixing the Broken Supply Chain

Each pillar includes multiple lines of effort and aggressive timelines for implementation. There are gaps—additional actions needed to prepare for an extended conflict requiring mobilization—but the starting point has been well articulated.

For example, in the first pillar, peacetime stabilized demand signals may suffice to stimulate private capital investment, but additional pre-arrangements are needed for government and industry investment in the capacity to meet wartime demands. The emphasis on commercial preference and industry-driven solutions will expand peacetime relationships with commercial firms whose capacity can be rapidly switched to defense production in time of war. The Industrial Base Consortium (part of Pillar 1) includes many non-traditional defense suppliers and thus can provide the DoW with insight into where capabilities and capacity for wartime exist in the industrial base.

Similarly, in Pillar 2, the focus on portfolio capability acquisition, management and trade space coupled with budget flexibility will transform peacetime response times in ways that can enable rapid wartime response. Again, additional steps are needed to pre-plan expanded capacity for replenishment of key warfighting capabilities.

In Pillar 3, digitizing acquisition can greatly increase speed of response, as discussed in the next section of this paper. Reforming Federal acquisition regulations and accounting requirements will address many of the long-standing barriers to integration of military and



commercial manufacturing. Additionally, modernizing test and evaluation and reforming accounting and acquisition regulations are essential to attracting commercial suppliers to government business.

Pillar 4 establishes the Wartime Production Unit (WPU), with the mission to transition the U.S. industrial base from a lean, “just-in-time” peacetime model to a “just-in-case” wartime footing capable of rapid surge and sustained high-intensity output. This is a key element of the strategy to prepare for industrial mobilization. The WPU uses a dedicated team of venture capital-style negotiators and manufacturing experts to negotiate multi-year deals with private companies to invest in capabilities and capacity needed in wartime. These deals are executed in close partnership with the DoW Economic Defense Unit (EDU)’s authorities for loans, grants and purchase commitments that incentivize private capital investments. Under the Arsenal of Freedom Framework, the WPU can license designs to multiple secondary manufacturers to ensure that if one factory is compromised or delayed, others can respond. Pillar 4 also emphasizes the use of modular open system architecture to increase the resilience of sources of supply and to provide for rapid upgrades in response to changing battlespace needs.

Pillar 5 addresses the need for technical data for weapon repair and maintenance. During acquisition, the DoW needs to balance industry intellectual property rights with government needs for depot repairs and competitive spare parts production. Historically, data delivered to the government has rapidly become outdated. New approaches such as Technical Data as a Service (Thompson & McGrath, 2019) offer ways to ensure digital data is up to date. Pillar 5 also addresses the need for supply chain visibility. Rather than a snapshot in time, the DoW will need to maintain visibility of the supply chain as it evolves over the lifecycle and continually mitigate supply chain risks.

In all of these pillars, the emphasis on commercial practices is vital to planning for large scale industrial mobilization. U.S. and allied commercial industries operate at a scale orders of magnitude larger than DoW peacetime demands. These industries are capable of delivering military products from commercial lines. This was demonstrated in an Air Force manufacturing technology program that used a commercial automotive production line to produce and qualify a circuit card for the F-22 aircraft (Air Force Research Laboratory, n.d.). The supplier, TRW, reported that the entire lifecycle requirement for this product, including production installs and lifetime spares, could be met by two shifts in their production plant. Wartime surge production will need to harness this commercial capacity. The National Academies report on Integrating Commercial and Military Manufacturing (National Research Council, 2002) provides additional examples and identifies barriers. Most of those barriers are now addressed by the Acquisition Transformation Strategy.

## **The Ecosystem: Effectiveness and Productivity Enhancers**

Through the DoW Acquisition Transformation Strategy, Government and Industry have begun to operationalize the ecosystem that will tap commercial production and manpower and provide digital transformation capabilities to reduce response times. Collaboration across an ecosystem of Government, Industry and Academic partners requires four key ecosystem elements to operate in concert with one another, as shown in

Table 4.

The first element, Strategic Governance, establishes clear authority, defines surge-readiness metrics, establishes the needed Government and Industry partnership structure, and authorizes “switch-on” contracting mechanisms for rapid scale-up across the broader Industrial Base to include commercial/non-traditional defense companies. The



Acquisition Transformation Strategy will provide solutions to address previous blockers, including:

- **Strategic Direction** – This includes the establishment of new governance bodies to drive more Government and Industry collaboration, including Business Operators for National Defense (BOND; Todd Lopez, 2026), Office of Industrial Base Policy (IBP) who help shape policy for streamlining acquisition, a Wartime Production Unit for accelerated munition production, and an Economic Defense Unit which acts as a financial accelerator for the defense industrial base (DoW, 2025c).
- **Organizational Execution** – The Strategy establishes new Portfolio Acquisition Executives (PAEs) as empowered, single-point accountability leaders responsible for consolidating and managing related weapon programs and the use of Adaptive Acquisition Framework Pathways that replace the old “one-size-fits-all” model.
- **Policy/Funding** – The Strategy leverages authorities and funding like those from the Defense Production Act (DPA), the newly established Joint Acceleration Reserve (JAR) (DoW, 2025b), and the Office of Strategic Capital (OSC) that give the ecosystem the legal and financial agility required to translate a real-time battlefield demand or sortie-mix change into an actionable procurement trigger.

*Table 4. Ecosystem Elements*

| <b>Ecosystem Element</b>      | <b>Core Role in the Kill-Chain → Supply-Chain Loop</b>                                                                                                                                                                              | <b>DoW Acquisition Transformation Strategy Pillars Supported</b>                                                                                                                                                                                                                             |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Strategic Governance</b>   | Sets authority, defines surge-readiness metrics, creates partnership structures (e.g., BOND, IMB, WPU, EDU) and authorizes “switch-on” contracts that turn battlefield demand into procurement actions.                             | <ul style="list-style-type: none"> <li>• Rebuild the Arsenal of Freedom</li> <li>• Maximize Acquisition Flexibility</li> <li>• Elevate and Empower the Acquisition Workforce (through policy &amp; funding levers)</li> </ul>                                                                |
| <b>Digital Infrastructure</b> | Provides a secure, federated data spine (TDaaS, software and data repositories, zero-trust APIs) that makes battlefield demand and supply-chain capacity machine-readable, continuously available, and protected.                   | <ul style="list-style-type: none"> <li>• Develop High-Performance Systems (model-based engineering, rapid data-driven decisions)</li> <li>• Improve Lifecycle Risk Management (real-time risk analytics)</li> <li>• Maximize Acquisition Flexibility (enables agile contracting)</li> </ul>  |
| <b>Industrial Base</b>        | Converts scientific discoveries into producible parts, maintains validated digital twins, supplies process knowledge and a trained workforce, and expands capacity through OEMs, MIIIs, Tier-2/3 suppliers, universities, and labs. | <ul style="list-style-type: none"> <li>• Rebuild the Arsenal of Freedom (surge-ready industrial base)</li> <li>• Elevate and Empower the Acquisition Workforce (training, up-skilling)</li> <li>• Develop High-Performance Systems (advanced-manufacturing, innovative materials)</li> </ul> |
| <b>Operational Users</b>      | Generates real-time kill-chain demand (damage reports, sortie-mix changes) via warfighters, JADC2 nodes, and forward logistics; consumes                                                                                            | <ul style="list-style-type: none"> <li>• Maximize Acquisition Flexibility (rapid demand-driven contract triggers)</li> <li>• Improve Lifecycle Risk Management (continuous feedback loop)</li> <li>• Elevate and Empower the Acquisition Workforce (hands-on operational insight)</li> </ul> |



|  |                                                             |  |
|--|-------------------------------------------------------------|--|
|  | delivered materiel and feeds feedback to the digital spine. |  |
|--|-------------------------------------------------------------|--|

The second element, Digital Infrastructure, embodies the DoW Digital Standards Strategy (Defense Standardization Program Office, 2026) definition of “a foundational architecture supporting the software tools, training, communication, and interoperability required to facilitate the exchange of real-time data across multiple platforms, organizations, and partners while maintaining strict access controls and cybersecurity.” To respond, the United States must realize an infrastructure that enables us to leverage all four Ecosystem Elements. Such an infrastructure must host a secure, federated data spine (e.g., software and data repositories, Technical-Data-as-a-Service [TDaaS], and zero-trust API gateways [National Security Agency, 2026], all in a secure, collaborative digital ecosystem) through which battlefield telemetry, requirement updates, and supply-chain capacity data become machine-readable, continuously available, and protected by strict access controls. This digital backbone can deliver the real-time demand signals the acquisition transformation pillars require, enable model-based engineering, and allow risk-based decision making across the entire product lifecycle at the speed of warfighter need.

The third element, Industrial Base (including academia), supplies the technical substance, workforce development, and private capital elements of the ecosystem. Systems Integrators, Original Equipment Manufacturers (OEMs), Manufacturing Innovation Institutes, mid-tier (Tier 2/3) suppliers, Universities, and National Laboratories convert scientific discoveries and technical innovations into manufacturable parts, maintain verified digital models and data for decision making, and populate the data lake with process knowledge, tooling footprints, and workforce competencies needed to realize the U.S. DoW vision. Furthermore, for the Industrial Base to be an effective deterrent, it must extend beyond conventional defense contractors to pre-plan and incentivize participation by non-traditional, commercial companies—a big area of emphasis in the Acquisition Transformation Strategy. Seth Jones (2025) underscores the importance of better leveraging the nation’s commercial sector not only to take advantage of technical advancements outside of defense, but also to extend capacity through small and medium size manufacturers. By integrating additional capabilities and capacity from the broader industry into a shared digital environment, the ecosystem fulfills the transformation goals of accelerating workforce up-skilling, expanding surge-capacity, and ensuring that high-performance systems can be prototyped, qualified, and fielded at the speed of relevance.

Finally, and most critical to ensuring clear articulation of demand from the battlefield, the ecosystem also must include the Operational Users themselves, including the warfighters, Joint All-Domain Command and Control (JADC2) nodes, and forward logistics operators. They provide the originating “kill-chain” inputs, consume the resulting materiel and capability produced by the Industrial Base to execute their mission and deliver battlefield effects. Shah and Kirchhoff (2024) highlight the importance of collaboration between military end-users and Industry developers from the earliest stages of development to iterate technology rapidly based on real-world feedback. Operational users’ real-time demand, damage reports, and logistics status close the feedback loop, supplying the data that drives the digital infrastructure and informs strategic governance decisions. When operational users can instantly query the collaborative digital ecosystem for part availability, supply chain surge-readiness levels, and delivery timelines, the acquisition transformation’s objectives of rapid flexibility and improved risk management become tangible.

To be prepared for a long duration conflict with a peer adversary, the ecosystem must evolve from a narrowly focused, short term surge model to an enabler of wartime





mobilization. That evolution requires: 1) expanding the industrial base to include commercial and nontraditional defense firms across the continental United States and in allied partner nations; 2) integrating these partners into a secure, standards based digital infrastructure that can convey real-time kill chain signals as actionable supply chain orders; and 3) embedding “switches” (e.g., preapproved surge readiness metrics, flexible funding authorities, and automated contract triggers) into the strategic governance framework of a secure and collaborative ecosystem. Only by deliberately and proactively identifying the actions needed for gap closure in each of the four ecosystem elements can the DoW and Industrial Base at large be ready to scale defense production to the 35–40% GDP range required for a full wartime economy.

Multiple recent workshops across the Aerospace and Defense Industry have shown that the creation of a Collaborative Open Digital Ecosystem (CODE) presents a unique opportunity for the DoW to fuse collaborative partnerships with a federated digital infrastructure, delivering a capability that can securely and directly link the kill chain to the supply chain at a tempo demanded by the modern warfare (American Institute for Aeronautics & Astronautics [AIAA] & NDIA, 2026). By securely linking battlefield capability demand data to a trusted network of government, industry, academia, and allied partners, CODE enables transformation of isolated acquisition silos into a dynamic, integrated “network of networks” required for realizing the U.S. DoW Warfighting Acquisition System (WAS). This collaborative ecosystem could provide the essential conduit through which rapid, data-driven decisions can be made, enabling surge production, logistics, and sustainment to be executed with the speed and scale of modern contested operations.

An integrated digital infrastructure provides the foundation for realization of this vision. Establishing a federated Technical Data as a Service (TDaaS) platform enables the ingestion and normalization of operational demand signals from Joint All Domain Command and Control (JADC2), translates a kill chain event (e.g., a damaged wing spar or a sortie mix shift) into an instantly observable surge request, and broadcasts it to authorized suppliers and surge ready facilities. Simultaneously, the same infrastructure delivers a real time view of supply chain capacity, tooling availability, and workforce readiness to acquisition officials and systems integrators. This bidirectional transparency allows the industrial base to size production lines, preposition tooling, and allocate skilled labor before a crisis escalates, thereby helping close the latency loop that hampers rapid fielding.

Realizing this vision, however, requires closing several critical gaps. First, a coherent data strategy to make data machine readable and machine interpretable is only partially defined and implemented, limiting seamless exchange across domains and ecosystem partners. Second, tool integration is fragmented with proprietary engineering suites and disparate versioning which hinders the rapid sharing of data, models, simulations and process workflows. Third, the overarching architecture lacks a universally adopted open standards framework, resulting in multiple reference models that impede cross program interoperability. Finally, trust in the underlying models (ensured through provenance metadata, continuous verification, and role-based access controls) has not yet been institutionalized across the broader ecosystem. Addressing these gaps through collaborative development of Government and Industry roadmaps, joint standards committees, and sustained investment in workforce credentialing will be essential for CODE to deliver the warfighter the decisive, surge ready supply chain capability and capacity the DoW Acquisition Transformation Strategy envisions.

## **A Framework for Industrial Surge Planning**

As shown in Figure 2, the Acquisition Transformation Strategy prepares the DoW for an initial conflict. Platforms and initial stocks of weapons and consumables can be deployed to support an operational plan. In some instances (e.g., munitions), ongoing production rates may



be sufficient to support replenishment needs after the initial 60–90 days. But for most commodities, a surge in production will be needed. The key idea in the framework for surge is to define pre-planned “switches” that can be activated at day 0 of the conflict to tap the increased industrial capacity needed for prolonged combat.




**Figure 2. Planning for Industrial Surge**

These switches need to be built into procurement contracts to provide pre-planned options and capacity for surge production rates. Beyond that, since the wartime demands may exceed the capacity of peacetime production facilities and supply chains, arrangements with additional suppliers are needed to prepare firms focused on commercial markets to be available for rapid switchover to defense production. The FY 2026 Defense Appropriation Act included \$177 million to establish the Civil Reserve Manufacturing Network (Committee on Appropriations, 2025), a voluntary network of “dual-use” commercial manufacturing firms. These firms would be required to maintain the capability to rapidly “flip the switch” from commercial production to military production during a national emergency or conflict.

The Defense Logistics Agency (DLA) has long maintained programs in peacetime to ensure availability of materiel needed in wartime. The National Defense Stockpile (NDS) program stores “strategic and critical materials” that the United States cannot produce enough of domestically but are essential for national defense. Since the Cold War, the NDS has shifted from just storing raw minerals to focusing on “buffer stocks” of semi-finished goods and chemicals to protect the Defense Industrial Base (DIB) from global supply chain shocks. The DLA Industrial Preparedness program, also known as the Warstopper program, is focused on items that are vital to the wartime mission but not generally required in peacetime (Reece, 2025). Both the stockpile program and the Warstopper program have arguably been underfunded in recent years. They are vital to preparing for a rapid transition to wartime industrial mobilization.

As previously discussed, planning for a whole-of-economy response to wartime demands will require use of DPA authorities. In addition to current uses of DPA authorities, pre-planning could benefit from the use of Voluntary Agreements, as proposed by McGinn (2025). Voluntary agreements and plans of action between government and industry can provide peacetime benefits to the firms involved as well as wartime industrial capacity for national emergencies.

An important factor in rapidly expanding supply chains is the availability of technical data for the items to be produced. In some cases, these data are intellectual property held by the prime contractor that would need to be licensed to additional suppliers. In other cases, the government holds license rights for use of the data and can disclose it to additional suppliers. Keeping the technical data up-to-date has been a past problem for the DoW. Research presented at past Naval Postgraduate School symposia has provided alternatives for patent pooling (Goertner et al., 2025) and for Technical Data as a Service (Thompson & McGrath, 2019). These alternatives may be on the critical path for industrial mobilization.

The office of the ASW (Industrial Base Policy) is developing plans for the Civil Reserve Manufacturing Network. This will be an important element of a written plan for industrial mobilization. It is unclear to the NDIA Surge Working Group who in government has overall responsibility for whole-of-economy mobilization planning.

A final element of this framework for industrial surge is recurring exercises and simulations to stress test supply chain responses to wartime scenario demands. Industry needs to be involved. Formation of industry boards to participate in classified and unclassified settings is needed for various industrial sectors that the DoW considers important in industrial mobilization. A modern version of the boards formed in WWII is needed, with provision for participation by allied industrial sources that the United States can count on in wartime.

## **Recommendations for DoW Consideration**

1. Establish continuous government & industry wargaming for quantifying demand. Assign responsibility for consolidating DoW estimates of demand profiles in an extended conflict based on annual Defense Guidance and operational plans. Share these estimates in a classified setting with cleared industry and get industry inputs on where capabilities and capacity exist for the needed commodities. Involve industry annually in wargames to provide feedback on constraints and timelines. This could include the establishment of a classified scenario-working group that leverages lessons learned from the unclassified structure demonstrated with the NDIA Ad Hoc Industrial Surge Working Group (with government lead, industry partners, and academic analysts) while operating under appropriate DoW security controls.
2. Establish Industrial Mobilization Boards for critical manufacturing sectors. Expand the industrial base bench to include commercial and non-traditional supply chain participants. Use the Civil Reserve Manufacturing Network (CRMN) as the initial basis for planning industrial mobilization. Coordinate the plan with relevant agencies to build support for a whole-of-government response. Implement contracts or voluntary agreements with commercial firms to prepare for switchover to defense manufacturing short notice. Build out the supply chain ecosystem across the continental United States and with our allies. Within the Acquisition Transformation Strategy, continue to emphasize commercial contracting and use of commercial solutions to open commercial mobilization options.
3. Build out the digital ecosystem and data required to rapidly connect kill chain demand to supply chain response. Accelerate digital transformation and improved management of the technical data needed for repairs and rapid expansion of supply chains. Leverage current Industry momentum to build out a Collaborative Open Digital Ecosystem and Technical Data as a Service (TDaaS) solutions for securely and rapidly sharing U.S. DoW and warfighter demand with the trusted industrial base and increase visibility of supply chain capability and capacity with the DoW and warfighter.



4. Work proactively with industry associations to inform industry of evolving plans and obtain broad industry input on how to align government and industry interests in preparing for a mobilization contingency. Develop a Surge Playbook as a national asset for how to rapidly pivot the broader industrial base (commercial and defense) from peacetime to wartime footing leveraging lessons learned on the needed switches and patterns established by the NDIA Ad Hoc Industrial Surge Working Group.

## Acknowledgements

The authors would like to specifically recognize the contributions of the NDIA Industrial Surge Working Group whose work significantly shaped this paper's content and conclusions. This Working Group included Major General Tim Zadalis, USAF (Ret.), Brench Boden, Rollie Dutton, and Valerie Sitterle. This core group, in partnership with the passionate workshop participants from across Government, Industry and Academia, was instrumental in informing this work.

## References

- Air Force Research Laboratory. (n.d.). (The defense industrial base: Integrated Product Team final report). 1999 (Report No. ADA367308).  
<https://apps.dtic.mil/sti/tr/pdf/ADA367308.pdf>
- Albon, C. (2026). *Space Command to launch wargame series for industry.*
- American Institute for Aeronautics & Astronautics & National Defense Industrial Association. (2026). *Toward realizing a collaborative open digital ecosystem* (AIAA SciTech 2026 Workshop Report [Unpublished], Issue NA).
- Cancian, M. F., Cancian, M., & Heginbotham, E. (2023). *The first battle of the next war.* Center for Strategic and International Studies.
- Cancian, M. F., Saxton, A., Bryan, L. A., & Helman, O. (2020). Industrial mobilization—Assessing surge capabilities, wartime risk, and system brittleness. *Naval Engineers Journal*, 132(2), 39–49.
- Cantrell, S. A., Margolis, C. H., Krauss, M. R., Pinon Fischer, O. J., & Mavris, D. N. (2025). *Digital twins for sustainment-oriented wargaming.* AIAA SciTech Forum, Orlando, FL.
- Crowley, D. (2022). *Engine fleet readiness services.* <https://standardaero.com/wp-content/uploads/2022/10/StandardAero-Maintenance-Insight.pdf>
- Defense Standardization Program Office. (2026). *Digital Standards Strategy.*  
[https://www.dsp.dla.mil/Portals/26/Documents/Publications/2026\\_Digital%20Standards%20Strategy\\_010826.pdf](https://www.dsp.dla.mil/Portals/26/Documents/Publications/2026_Digital%20Standards%20Strategy_010826.pdf)
- Federal Reserve Bank of Saint Louis (FRED). (2026). *Shares of gross domestic product: Government consumption expenditures and gross investment: Federal: National defense.*
- Goertner, F., Lucyshyn, W., Crocker, J., O'Brien, T., Bailey, J., Elmaghraby, W., Hill, H., & Huddleston, R. (2025). *Time value of data decision modeling for major defense acquisition programs.*
- Gunziger, M. A. (2021). *Affordable mass: The need for a cost-effective PGM mix for great power conflict.* Mitchell Institute for Aerospace Studies.
- Jones, S. G. (2023). *Empty bins in a wartime environment: The challenge to the US defense industrial base.* Bloomsbury Publishing.
- Jones, S. G. (2025). *The American edge: The military tech nexus and the sources of great power dominance.* Oxford University Press.
- McGinn, J. (2020). *Building resilience: Mobilizing the defense industrial base in an era of great-power competition.* The Heritage Foundation.





- McGinn, J. (2025). *Options for strengthening the use of Defense Production Act Title VII*. <https://acqirc.org/publications/research/innovation/options-for-strengthening-the-use-of-defense-production-act-title-vii/>
- National Defense Industrial Association Manufacturing Division. (2025). *NDIA & DRIVE contested logistics: Workshops summary & recommendations*. [https://www.ndia.org/-/media/sites/ndia/divisions/manufacturing/documents/ndia\\_drive\\_workshop-report\\_v5.pdf](https://www.ndia.org/-/media/sites/ndia/divisions/manufacturing/documents/ndia_drive_workshop-report_v5.pdf)
- National Research Council. (2002). *Equipping tomorrow's military force: Integration of commercial and military manufacturing in 2010 and beyond*. The National Academies Press. <https://doi.org/https://doi.org/10.17226/10336>
- National Security Agency. (2026). *Zero trust implementation guideline primer* (Cybersecurity Technical Report Issue). [https://media.defense.gov/2026/Jan/08/2003852320/-1/-1/0/CTR\\_ZERO\\_TRUST\\_IMPLEMENTATION\\_GUIDELINE\\_PRIMER.PDF](https://media.defense.gov/2026/Jan/08/2003852320/-1/-1/0/CTR_ZERO_TRUST_IMPLEMENTATION_GUIDELINE_PRIMER.PDF)
- O'Hanlon, M. E. (2024). *U.S. defense spending in historical and international context*.
- Reece, B. (2025). *Warstopper program changes could improve surge capabilities for critical supplies*. <https://www.dla.mil/About-DLA/News/News-Article-View/Article/4157222/warstopper-program-changes-could-improve-surge-capabilities-for-critical-suppli/>
- Roosevelt, F. D. (1940). *Message to Congress on appropriations for national defense*. <https://www.presidency.ucsb.edu/documents/message-congress-appropriations-for-national-defense-0>
- Shah, R. M., & Kirchhoff, C. (2024). *Unit X: How the Pentagon and Silicon Valley are transforming the future of war*. Simon and Schuster.
- Stockholm International Peace Research Institute. (2025). *Unprecedented rise in global military expenditure as European and Middle East spending surges*.
- Thompson, G. E., & McGrath, M. (2019). *Technical data as a service (TDaaS) and the valuation of data options*.
- Tian, N., Lopes da Silva, D., Liang, X., Scarazzato, L., Béraud-Sudreau, L., & Assis, A. (2024). *Trends in world military expenditure*. [https://www.sipri.org/sites/default/files/2025-04/2504\\_fs\\_milex\\_2024.pdf](https://www.sipri.org/sites/default/files/2025-04/2504_fs_milex_2024.pdf)
- Todd Lopez, C. (2026). *Department seeks counsel of industry leaders to advance Arsenal of Freedom*. <https://www.war.gov/News/News-Stories/Article/Article/4405845/departments-seeks-counsel-of-industry-leaders-to-advance-arsenal-of-freedom/>
- U.S. Department of War. (2025a). *Acquisition Transformation Strategy*. <https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/ACQUISITION-TRANSFORMATION-STRATEGY.PDF>
- U.S. Department of War. (2025b). *Reforming the joint requirements process to accelerate fielding of warfighting capabilities*. <https://media.defense.gov/2025/Nov/10/2003819442/-1/-1/1/REFORMING-THE-JOINT-REQUIREMENTS-PROCESS-TO-ACCELERATE-FIELDING-OF-WARFIGHTING-CAPABILITIES.PDF>
- U.S. Department of War. (2025c). *Transforming the Defense Acquisition System into the Warfighting Acquisition System to accelerate fielding of urgently needed capabilities to our warriors*. <https://media.defense.gov/2025/Nov/10/2003819439/-1/-1/1/TRANSFORMING-THE-DEFENSE-ACQUISITION-SYSTEM-INTO-THE-WARFIGHTING-ACQUISITION-SYSTEM-TO-ACCELERATE-FIELDING-OF-URGENTLY-NEEDED-CAPABILITIES-TO-OUR-WARRIORS.PDF>
- U.S. Department of War. (2026). *Department of War establishes new acquisition model to more than triple PAC-3 MSE production in partnership with Lockheed Martin*. <https://www.war.gov/News/Releases/Release/Article/4371320/departments-of-war-establishes-new-acquisition-model-to-more-than-triple-pac-3-m/>





Vergun, D. (2020). During WWII, industries transitioned from peacetime to wartime production. *U.S. Department of Defense*, 27, 233–254.

Welch, K. (2026). *Wargames keep warning us about congested logistics—It's time to take action*. <https://mwi.westpoint.edu/wargames-keep-warning-us-about-congested-logistics-its-time-to-take-action/>



# Establishing a Manufacturing Security Program

**Dr. Jerry McGinn**—Director of the Center for the Industrial Base and senior fellow with the Defense and Security Department at the Center for Strategic and International Studies (CSIS) in Washington, D.C.

**Katy Buda**—Associate director at the Center for the Industrial Base.

**Celia Barrie**—Former interns with the Center for the Industrial Base.

**Sarah O'Rourke**—Former interns with the Center for the Industrial Base.

## Abstract

The Department of Defense faces persistent challenges related to munitions production, including single source suppliers, reliance on foreign-sourced materials, limited surge capacity, and cybersecurity threats to defense industrial systems. These vulnerabilities compromise readiness and strategic deterrence while endangering the warfighter. This paper provides foundational research for the concept of a Manufacturing Security Program (ManuSP), a framework designed to ensure secure, scalable, and resilient domestic manufacturing of critical munitions. The paper analyzes seven case studies for their effectiveness in the following four components: 1) authorities and policies, 2) government-industry collaboration; 3) resources and contracting 4) and activation phases and utilization. The case studies focus on the following programs: the Voluntary Intermodal Sealift Agreement (and Maritime Security Program), the Voluntary Tanker Agreement (and Tanker Security Program), the "Voluntary Agreement, Manufacture and Distribution of Critical Healthcare Resources Necessary to Respond to a Pandemic," the Civil Reserve Air Fleet, the National Armaments Consortium, the Munitions Campus, and the Civil Reserve Manufacturing Network. They are scored as either "low," "medium," or "highly" effective. The following paper is a draft of the ongoing research project.

## Voluntary Intermodal Sealift Agreement Case Study

### Overview

The Voluntary Intermodal Sealift Agreement (VISA) is a Defense Production Act (DPA) Title VII voluntary agreement established to provide the Department of Defense (DoD) with guaranteed access to commercial vessels for sealift and intermodal shipping services and systems during times of conflict or emergency. VISA also provides participants with increased peacetime business contracts and opportunities (Maritime Administration, 2024). Of these vessels, a subset are part of the Maritime Security Program (MSP). MSP, which shortly predated VISA, uses financial incentives to encourage carriers to participate in VISA to grow participation in the VISA fleet (McGinn & Letts, 2025). MSP also focuses on sustaining a labor resource of American mariners and enables operators to make upgrades to their vessels (Maritime Administration, n.d).

### Authorities and Policies

VISA falls under Section 708 of Title VII of the DPA and first went into effect in 1997. Section 708 enables the President to work with representatives of industry, finance, business, agriculture, labor, and other sectors to create voluntary emergency preparedness agreements in the instance that there is a direct threat to national defense or emergency preparedness programs (Maritime Administration, n.d.). As a voluntary agreement, VISA must be approved by the Department of Justice (DoJ) and Federal Trade Commission (FTC) to allow for the exemption of antitrust laws. This process is difficult, and there are arguments that voluntary agreements should only be created where other options for government-industry collaboration don't exist (McGinn & Letts, 2025).

The driving force for establishing MSP and VISA was preventing the reduction of the U.S.-flag fleet and better maintaining trained personnel for U.S. shipping and transportation



needs. VISA took at least 6 years to negotiate after the Gulf War and was created in part because of the failings of the Sealift Readiness Program (SRP). SRP was not used during Operation Desert Storm or Operation Desert Shield because it was not responsive enough during wartime and caused financial risk for SRP companies (McGinn & Letts, 2025). At the end of the 20th century, MSP saved the two largest American commercial carriers, American President Lines and Sealand, from having to transition vessels to foreign flags. The U.S. Ocean Shipping Reform Act of 1998 allowed for confidential services contracts and was instrumental in helping U.S. commercial carriers negotiate with the U.S. government without needing to be the lowest bid (Lewis & Coulter, 2000). Since VISA is a preexisting program, it is easier to acquire a renewal, which takes between 6 to 8 months, but if a new voluntary agreement were established, it would be more difficult due to the need for an antitrust exemption. For renewals, the Maritime Administration (MARAD) must make a case why VISA is achieving its purpose and then the DoJ Antitrust Division and the FTC must sign off on the renewal (McGinn & Letts, 2025).

MARAD, under the Department of Transportation, is the sponsorship authority for VISA, and the United States Transportation Command (USTRANSCOM) jointly maintains the VISA fleet. USTRANSCOM, MARAD, and the participating commercial shipping companies work together to determine the available capacity of participants to supplement the DoD at different stages of a conflict or emergency. Carrier Coordination Agreements (CCAs) provide flexibility to participants so that they can navigate commercial operation needs with DoD contingency efforts. The federal government also has the option to requisition non-VISA ships in times of emergency through Title 46, but the purpose of VISA is to prepare ahead of time with a framework in place to support the DoD. USTRANSCOM is the procurement authority for ships under VISA, while USTRANSCOM and MARAD then maintain and operate the VISA fleet (Maritime Administration, n.d.).

The Maritime Security Act of 1996 created the MSP (Maritime Administration, 2005). The 2020 National Defense Authorization Act (NDAA) extended MSP operating agreements until September 30, 2035 (Maritime Administration, 2025). MARAD can maintain at most 60 vessels in the MSP program in a single fiscal year (FY; Maritime Administration, 2023). MSP is a financial incentive for industry participation. All MSP carriers must be a part of VISA, which is an emergency preparedness program, but not all VISA carriers are a part of MSP. MSP carriers are required to commit 100% of their MSP capacity and associated intermodal transportation resources to the third activation stage of VISA, while the remaining VISA operators must commit only 50% of their U.S.-flag capacity (Maritime Administration, n.d.). The goal of MSP payments is to help companies counter the cost of running U.S. flag vessels which can be more costly than foreign flag vessels. MSP agreements are established per company, not vessel (Chatham House Interview, 2025).

VISA and MSP promote yearly government-industry collaboration and offer clear forums and open channels for communication. MARAD and USTRANSCOM co-chair the Joint Planning Advisory Group (JPAG) which allows their organizations and VISA participants to coordinate (Maritime Administration, 2024). JPAG is a planning forum for reviewing DoD sealift and intermodal contingency requirements, analyzing commercial capacity available to meet those requirements, and creating Concepts of Operation (CONOPS) for such requirements (Maritime Administration, 2024). The strength of authorities and policies for VISA rests on its ability to be renewed relatively quickly and the clear structures and legislation in place. **Assessment: High**

### **Government-Industry Collaboration**

There is a clear and straightforward financial incentive for industry to take part in MSP and VISA. For VISA vessels, they receive priority access to DoD cargo when the United States is not at war. This access has been explained as “peacetime business for wartime capacity.”



VISA commercial carriers have transported cargo for the United States to both Afghanistan and Ukraine where at one time about 85% of sustainment support was provided by VISA or Civilian Reserve Air Fleet participants (McGinn & Letts, 2025). MSP participants benefit further by receiving funding per ship. For MSP, the federal government pays MSP participants based on available appropriations and specific funding ceilings. The roster of vessels in the MSP can change, and new applications are accepted through the Federal Register when there are openings for new operating agreements (Maritime Administration, 2023). One of the benefits of the VISA program is that the program does not need to be activated for there to be productive direct negotiations and contracts between the federal government and industry. There are opportunities for government-industry engagement through the two unclassified VISA meetings that occur every year between the government and carrier companies. These meetings provide an avenue for open discussion of industry capacity and the tradeoff of risks and opportunities (McGinn & Letts, 2025). One of the strengths of voluntary agreements and VISA is the ability for government-industry collaboration, even without the activation of the agreement. **Assessment: High**

### **Resources and Contracting**

Congress allocates funding in the U.S. Department of Transportation's budget for the MSP (Maritime Administrator, n.d.). Congress authorizes and funds the program in increments, so continued funding would depend on the priorities of Congress over time, but there is an established process of authorizing and providing funding for MSP. Personnel for the program depend on staffing at MARAD and USTRANSCOM. As both organizations have functions beyond the MSP and were not stood up solely to execute the program, there is greater stability in continued staffing and oversight for MSP. MARAD issues awards for new MSP Operating Agreements (in keeping with the 60-vessel maximum for the fleet), and MARAD then places a notice for applications in the Federal Register, highlighting which vessels are needed. MARAD reviews the applications and associated vessels to determine their commercial viability while USTRANSCOM assesses their potential use for the military. When both MARAD and USTRANSCOM agree on a vessel, MARAD awards the new Operating Agreement to the chosen carrier (Maritime Administrator, n.d.).

The antitrust exemption means that MARAD can directly negotiate with VISA participants regardless of activation which facilitates open communication, discussions of resource and capacity gaps, and enables better logistics tracking for the U.S. government. Private companies can better track their long-term business opportunities (McGinn & Letts, 2025). VISA is a form of soft requisition that is supported by established government-industry relationships. During conflict, the government does not need to use open bids and can bypass the typical 120-day contracting cycles (Chatham House Interview, 2025). The federal government does not need to use an acquisition pathway as they are not purchasing new ships for MSP but only creating Operating Agreements to use commercial vessels when necessary. There is an established process for these agreements, and the program has been running for several decades, so the contracting process is relatively straightforward. **Assessment: High**

### **Activation Phases and Utilization**

There are three activation stages to the VISA program. Stage I and II allow for contingency contracts between USTRANSCOM and the participating companies to meet expected DoD Contingency requirements. The contracts are negotiated ahead of time, and the agreements are carried out in line with DoD contracting methodologies. Stage III offers the DoD more capacity if Stages I and II capacity cannot fulfill Contingency requirements and there aren't enough shipping services available from non-VISA participants through standard DoD contracting practices or U.S. government treaties. The Secretary of Defense is able to activate each stage of VISA, and, when required, MARAD will report to the FEMA Administrator on the



state of the agreement. There is also a mandatory position reporting system for vessels in VISA (Maritime Administration, 2024).

The main issue in assessing the activation phases is that VISA has never been activated, but the benefit of close ties between industry and government through VISA means that activation would likely be simple and effective (Kok, 2021). McGinn and Letts explain “voluntary agreements have saved money and supported national security objectives through their collaborative function in peacetime scenarios” (McGinn & Letts, 2025). Activation would most likely be an effective process, but there have been no activations, and therefore MARAD and TRANSCOM have no experience to draw on. **Assessment: Medium**

## **Voluntary Tanker Agreement**

### **Overview**

The Voluntary Tanker Agreement (VTA) is a voluntary agreement through DPA Title VII that enables the U.S. government to contract commercial tankers during national emergencies. The agreement is designed to ensure the DoD can draw on commercial product tanker capacity to support national defense and global operations (Maritime Administrator, 2022a).

### **Authorities and Policies**

As with VISA, Section 708 of the DPA authorized the establishment of the VTA. Originating in 1951, the VTA is one of the oldest voluntary agreements. Activated to provide support to meet fuel transportation needs during the Korean War, the VTA remained on standby until it lapsed in 2013. The program halted largely due to limited industry participation, driven by lackluster government incentives. Due to increased concerns over national security and emergency response capabilities, the U.S. government decided to reactivate the VTA (McGinn & Letts, 2025). In addition to the VTA, MARAD the Tanker Security Program (TSP), which provides financial incentives for companies to enroll in the VTA. The TSP strengthens industry’s ability to participate in the international trade of liquid fuel products and the DoD’s capability to respond to emergencies (Maritime Administrator, 2022b).

As with the VISA program, USTRANSCOM and MARAD oversee the VTA fleet. MARAD serves as the program’s primary oversight authority and the DoD retains authority over program activation. USTRANSCOM is responsible for the procurement of commercial tankers and manages the contracts that would be used during activation. MARAD has general oversight and authority power over the obligated tanker fleet (Maritime Administrator, 2022a). The VTA has two annual meetings with the tanker and carrier companies to align priorities, however if activated, the VTA would formally convene to establish a joint planning advisory group (McGinn & Letts, 2025).

Ultimately, the goal of payments through the TSP is to ease the financial burden of companies using U.S. flag vessels and encourage their use (Chatham House Interview, 2025). There are well-established government authorities and clear policies for industry participants. The VTA’s long-standing history and its successful reactivation demonstrate the program’s ability to operate effectively within government processes, including coordination with the DoJ and FTC to establish and reaffirm often complicated and lengthy legal agreements.

**Assessment: Medium**

### **Government-Industry Collaboration**

The VTA provides industry participants with both financial and legal incentives. This program was designed to create a close working relationship between industry partners, MARAD, and USTRANSCOM. In addition to receiving funding in exchange for each participating tankers, companies enrolled in the VTA will receive priority contracting for





peacetime operations (McGinn & Letts, 2025). The DoD provides funding for any increased war risk insurance premiums for vessels. Furthermore, the VTA provides participants with civil and criminal defense for potential violations of antitrust laws while carrying out VTA contractual obligations. Tankers participating in the VTA remain on continuous standby to support DoD operational needs. The VTA provides an avenue for direct negotiations between government and industry, having created well-integrated forums for participation that meet twice a year. These meetings allow industry professionals and government officials to collaborate and communicate about tanker capabilities, program priorities, and industry capacity. It also provides an outlet for the two groups to come together and discuss their risk assessments, allowing them to address and mitigate concerns both parties may have (Maritime Administration, 2022a). While there are clearly articulated contracts and incentives for industry partners, these incentives have historically proven to be too lackluster and the program is at risk of not providing a continuing incentive structure to support longevity. **Assessment: Medium**

### **Resources and Contracting**

Following its reauthorization, the VTA received funding through congressional appropriations for TSP. Section 3511 of the FY21 NDAA approved funding for and authorized the TSP. Congress then appropriated \$60,000,000 in funding for the TSP in the FY22 Consolidated Appropriations Act. The funding provides \$6 million per tanker, per year for up to 10 U.S. Flag tankers. This funding is to remain until it is depleted, at which point new funds will be requested and need to be reauthorized and appropriated through Congress. Participating vessels are required to meet specific standards, undergo inspections, and maintain necessary capabilities to uphold DoD requirements and meet MARAD safety standards. The vessels must maintain a Certificate of Inspection from the U.S. Coast Guard to continue participation in the program. Commercial participants are responsible for the maintenance, overhaul, and repair work necessary to meet program requirements. The companies use funding they receive from the government to offset these costs (Maritime Administrator, 2022b). The VTA represents a form of soft requisition where the government does not need to use open bids. When a conflict occurs, the VTA enables the government to pursue quick contracts without the typical 120-day cycle. Relationships with industry are also established in advance and help streamline contracting during conflicts (Chatham House Interview, 2025). **Assessment: Medium**

### **Activation Phases and Utilization**

The VTA can be activated by the Commander of USTRANSCOM, under the direction of the Secretary of Defense. Activation is approved by the Secretary of Defense when it is sufficiently demonstrated that the existing tanker capabilities are insufficient to respond to the national emergency. Following the approval of the activation, MARAD is notified, which in turn notifies the FTC Chair and the Attorney General. When the VTA is activated, they do not ask each company to provide a certain amount of space on the tankers but instead contract the tankers in their entirety. The number of tankers contracted is dependent upon the need of the government at the time. As of this writing, the VTA has not been activated since its original activation during the Korean War. VTA activation uses participant tanker capacity in timed phases. There is a priority for capacity, with the stages as follow: 1) U.S.-Flag vessel capacity operated by a VTA participant, 2) U.S.-Flag vessel capacity operated by a carrier that is not a participant, 3) combination of U.S./foreign flag vessel capacity operated through a VTA participation, 4) combination of U.S. owned or operated foreign flag capacity operated by a non-participant, and lastly 5) foreign-owned or operated foreign flag capacity operated by a carrier outside the VTA program (Maritime Administration, 2022a). This category has not been activated in over 75 years, and uncertainties remain regarding how the process would function in practice. Potential challenges may not become apparent until an activation occurs.

**Assessment: Low**



## COVID-19 Case Study

### Overview

The COVID-19 pandemic revealed severe shortcomings in the U.S. industrial base's capacity to meet domestic medical supply needs, largely due to the nation's heavy reliance on foreign sources. By leveraging Title VII of the DPA, the federal government created the Voluntary Agreement for the Manufacture and Distribution of Critical Healthcare Resources Necessary to Respond to a Pandemic. This Covid voluntary agreement united private-sector participants and the Federal Government to expand the manufacturing and distribution of critical healthcare resources nationwide for pandemic response through coordinated planning, information sharing, and collaboration in support of national defense needs. However, this agreement was never fully utilized because it was executed too late. The Covid voluntary agreement concluded in August 2025 (FEMA, 2021).

### Authorities and Policies

The Covid voluntary agreement was designed to expand the effectiveness of manufacturing and distribution of healthcare supplies by uniting the efforts of private-sector participants and the federal government. It was first authorized in August 2020. In President Trump's Executive Order (E.O.) 13911, the Secretary of Health and Human Services and the Secretary of Homeland Security were given the authority of the President per Section 708(c)(1) and (d) of the DPA to enact the Voluntary Agreement (Trump, 2020). The Federal Emergency Management Agency (FEMA) was designated as the Sponsor for this agreement under Section 708 of the DPA (FEMA, 2021).

As part of the DPA-required consultation process, FEMA, as the sponsoring agency, coordinated closely with several key federal partners when forming the Voluntary Agreement. These consulting federal agencies included the Department of Health and Human Services (HHS), the DoJ, and the FTC, including its Chair (FEMA, 2021). This interagency coordination helped ensure that the agreement aligned with public health priorities, legal requirements, and anti-trust safeguards while supporting the broader national defense mission (Department of Homeland Security, 2020). However, due to minimal waiting periods, such as the approval of the Attorney General, fully authorizing the Agreement took longer than expected (GAO, 2025).

In addition to federal oversight, the organizational structure of the agreement was built around private-sector participation. Private companies, manufacturers, suppliers, distributors, trade associations, and other non-federal organizations could join as "Participants" once invited by FEMA and upon meeting the eligibility criteria. The agreement included six different plans of actions, defined as "documented methods adopted by participants in an existing voluntary agreement to implement that agreement" (FEMA, 2021). After signing a Plan of Action Agreement, participants committed to collaborative efforts under the Voluntary Agreement, including coordinated planning, manufacturing, information sharing, allocation, and distribution of critical healthcare resources to support pandemic response efforts. To implement each Plan of Action, FEMA could convene structured meetings and established specialized subcommittees between industry and government focused on specific supply categories or logistical challenges, such as respirators, gloves, gowns, or eye protection (FEMA, 2021).

The 2020 Voluntary Agreement was designed to coordinate the manufacture and distribution of critical healthcare resources only within the United States for pandemic response (FEMA, 2021). Its language and implementation focus on nationwide production, allocation, and domestic supply-chain coordination through FEMA-led plans of Action. The authorities of this framework were defined, however, the delay in establishment made the Agreement ineffective. A former government representative who was responsible for driving the formation of the Covid Voluntary Agreement explained that government lawyers and personnel lacked a foundational



knowledge of voluntary agreements. Another temporal barrier was the requirement that each step goes through the Attorney General (Chatham House Interview, 2025). **Assessment: Low**

### **Government-Industry Collaboration**

Private industry played a large role in establishing this voluntary agreement. The Agreement allowed FEMA, after consultation with HHS, DoJ, and FTC, to convene with private-sector suppliers (81st Congress, 1950). Subcommittees brought together federal officials and private-sector participants to assess demand, identify production bottlenecks, coordinate sourcing of raw materials, and align distribution with national needs (FEMA, 2021). Meetings held under the authority of Section 708 served as the primary mechanism for guiding implementation, and expanding participation as needed (FEMA, 2021). For the Covid Voluntary Agreement, morning meetings were run with primarily eight companies, though a far larger list of companies signed up for the Voluntary Agreement. It is estimated that approximately 80% of companies that joined the Covid voluntary agreement were not relevant to the pandemic response effort (Chatham House Interview, 2025).

Implementation was designed to occur through resource-specific Plans of Action, each tailored to a particular category of healthcare resources. The first such plan, the Plan of Action to Establish a National Strategy for the Manufacture, Allocation, and Distribution of Personal Protective Equipment (PPE), created multiple subcommittees, bringing together government representatives and industry participants to ensure the timely production and distribution of PPE across the country, manage the nation's stockpile of PPE, establish a streamlined process for FEMA Allocation of PPE, and ensure competition in manufacturing of PPE under the DPA (Department of Homeland Security, 2020). The other five plans followed the same framework, aimed at tackling manufacturing shortages for healthcare supplies. Participants in a Voluntary Agreement have antitrust protections and were given defenses to civil and criminal actions brought for violations of antitrust laws. However, this protection only applied while the agreement was authorized (FEMA, 2021).

The Agreement created a structured process for accelerating the production of supplies by intersecting private-sector capabilities with federal priorities. Antitrust protections incentivized firms to collaborate and scale manufacturing quickly, while Plans of Action established production targets and allocation strategies (FEMA, 2021). Unfortunately, the COVID Voluntary Agreement was established too late and was not in place to prevent the race to secure PPE and the eventual scarcity of such necessary protection equipment. The Agreement could have been a key channel for industry and government to collaborate in maximizing existing sources of PPE. Furthermore, with a voluntary agreement established earlier in the pandemic, industry might have had more time to respond by onshoring production, increasing the domestic prevalence of such material production for future need (Chatham House Interview, 2025).

**Assessment: Medium**

### **Resources and Contracting**

The Coronavirus Aid, Relief, and Economic Security (CARES) Act was enacted in March 2020 to provide broad and economic and public health support. Contributing nearly \$2.2 trillion, the CARES act provided agencies such as FEMA, DoJ, and HHS with funds to carry out tasks related to pandemic support (Peter G. Peterson Foundation, n.d.). FEMA struggled to establish the Agreement in time to have a meaningful effect on the production of healthcare supplies, in part, because of minimum waiting periods (GAO, 2025). Moreover, its effectiveness was limited by participants' ability to adjust operations dynamically and supply-chain constraints (GAO, 2021). **Assessment: Low**



## Activation Phases and Utilization

Implementation of the Covid Voluntary Agreement was carried out through formal, resource-specific Plans of Action approved by FEMA. The first and most significant Plan of Action focused on establishing a national strategy for the manufacture, allocation, and distribution of PPE in response to COVID-19. The other five Plans focused on diagnostic test kits, drug products and substances, medical devices, and supply chain logistics. Once a Plan of Action was approved, participating companies signed onto that specific plan, which defined the scope of permitted coordination, the objectives of production and distribution, and the safeguards governing information sharing and collaboration (FEMA, 2021). Unfortunately, the Covid voluntary agreement was designed and implemented too late to be properly activated and effective. **Assessment: Low**

## Civil Reserve Air Fleet

### Overview

The Civil Reserve Air Fleet (CRAF) is a voluntary partnership through which U.S. commercial air carriers provide surge capacity to the United States military during national emergencies in exchange for priority access to government business in peacetime. CRAF demonstrates the ability to augment U.S. military resources through the mobilization of existing domestic and commercial resources. All major U.S. airlines, along with many regional carriers and cargo companies, have been consistent participants in CRAF (Department of Transportation [DOT], 2024).

### Authorities and Policies

The DPA established broad presidential authorities to enable the mobilization of the U.S. industrial base in support of national security (Neenan, 2023). Title I of the DPA empowers the executive branch to incentivize commercial businesses to prioritize and accept government contracts that service national defense and sustain warfighting capabilities. Established in 1951 as a joint initiative of the DoD and the Department of Commerce, CRAF was stood up one year after the creation of the DPA, as one of the original DPA initiatives (Russell, 2025). CRAF has had limited structural changes over its lifespan. The most significant change in program history was the transition of joint authority between the DoD and the Department of Commerce to the DoD and the Department of Transportation (DOT; DOT, 2024). As a joint program, the operations of CRAF are split between the DOT and DoD; the DOT creates and manages contracts with commercial passenger airlines and cargo companies, while USTRANSCOM manages the activation phases (Peters, 2021). The government enacts thousands of prioritization actions a year which means the government often “jumps itself in line” (Chatham House Interview, 2025).

CRAF participants are required to meet performance and regulatory standards to ensure that if they are activated, they can complete assigned tasks to expectations (DOT, 2024). This includes requiring airline personnel involved in CRAF to maintain security clearances, maintain strict compliance with FAA regulations, meet performance and safety standards, be able to commit 40% of their CRAF-capable fleet, and maintain four complete crews per aircraft (Peters, 2021). CRAF is organized into two primary participation segments—national and international—each designed to support distinct operational requirements. The international segment provides overseas airlift support, while the national segment supports domestic and intra-theater missions (Peters, 2021). Through CRAF, the federal government can indemnify companies and protect them from being sued when they must change commercial flights and prioritize government needs (Chatham House Interview, 2025). Executive Order 13603 (National Defense Resource Preparedness) signed by President Obama on March 16, 2012, reinforces that the activation of CRAF is executed when determined necessary and under the discretion of





the Secretary of Defense (Obama, 2012). CRAF's authorities and policies are well established, clearly articulated, and supported by defined oversight mechanisms. **Assessment: High**

### **Government and Industry Collaboration**

Under CRAF, the government provides participating companies with priority access to government business. The U.S. government is highly reliant on domestic passenger and cargo carriers to transport its goods, mail and people which provides robust business for CRAF participants (Baldanza, 2021). Additionally, CRAF participants have the potential to receive tax incentives and are eligible for broader U.S. government transportation contracts (Warren & Conklin, 2024). When CRAF is activated, the government pays the carriers for their services at a pre-determined rate that covers both the cost and provides a return on investment (Lyman, 2011). The CRAF model requires industry to provide the government with detailed financial information and through close coordination and communication, CRAF has promoted greater understanding between government and industry while providing useful relationships and contacts for USTRANSCOM (Chatham House Interview, 2025).

Major passenger and cargo carriers, and many smaller airlines, have remained committed participants in both the international and national segments of CRAF largely because of the strong incentives for industry. The properly structured firm-level incentives and the low likelihood of mobilization underpin CRAF's durability and strength over time. The operational demands placed on activated CRAF participants closely mirror their normal commercial activities. Government and industry collaboration through CRAF is successful because there are clear and well-established incentives for industry and the government has proven to be a reliable partner—both in providing business and limiting mobilization until necessary (Harshberger, n.d.). **Assessment: High**

### **Resources and Contracting**

CRAF participation is renewed on an annual basis, with contracts and funding reflecting the needs of the Air Force and composition of the fleet and business needs of the industry. The Mobility Capabilities and Requirements Study (MCRS), a classified study conducted by the DoD every five years, determines the necessary capacity and makeup of CRAF. This analysis guides the strategic planning in the buildout of CRAF agreements and the recruitment of industry participants (Tirpak, 2024). As of 2010, the DoD obligated \$26.20 billion to various CRAF participants, reflecting the program's sustained operational value to U.S. mobility planning. In FY2025, USTRANSCOM obligated \$1.8 billion for CRAF passenger and cargo airlift services (Federal Procurement Data System, 2025). Congressional funding must be renewed each fiscal year, and lapses or delays in the budget process could complicate the negotiation of new agreements or interfere with the activation of current participants if necessary. CRAF is also considered an expensive program (Chatham House Interview, 2025). That said, there have been no major issues maintaining the program historically, and CRAF is popular in Congress.

**Assessment: High**

### **Activation Phase and Utilization**

Since its inception, CRAF has been activated on three occasions; Operation Desert Storm/Shield (1990–1991), Operation Iraqi Freedom (2002–2003), and Afghanistan Noncombatant Evacuation Operations (2021). Each of these activations were classified as stage one activations. Stage one is activated for “humanitarian assistance, disaster relief, and minor regional crises.” Stage one is activated largely during peacetime or in a minor regional crisis and is called upon when AMC airlift capability cannot meet deployment expectations and other airlift requirements simultaneously. This stage exclusively relies on international, long-range aircraft. Stage two activations are for “major theater wars.” This activation is used to support defense emergencies that are greater than stage one but do not yet meet the





requirements of a full national emergency. This stage relies on all three segments: national, international long-range, and international short-range. Stage three is activated during “periods of national mobilization” (DoD, 2021). Stage three activations mobilize the entirety of the CRAF fleet. This can occur when there are either, multiple major theater wars, a national emergency declared by the President, or Congress or a serious national security situation. CRAF can be activated by the Commander of USTRANSCOM with approval from the Secretary of Defense in stage one and two. Stage three requires authorization from the Secretary of Transportation as well as the Secretary of Defense. It is expected that when mobilized, CRAF carriers will have their aircraft ready within 24–48 hours after activation and mission assignment. The Air Mobility Command assigns missions dependent upon the stage of activation. When CRAF is deactivated, it is required that the AMC provide the carriers with a 15-day notice (Lyman, 2011).

The emergency evacuation from Afghanistan demonstrates a successful CRAF activation. In 2021, Secretary of Defense Austin III ordered the stage one activation of CRAF to augment the Department of State’s efforts to evacuate U.S. citizens and personnel and at-risk individuals from Afghanistan. Through this activation, 18 aircraft from United, American, Delta, Omni, Hawaiian, and Atlas Air were sent to interim staging areas to move people to their final destinations. The activated aircraft did not fly directly into the crisis area in Kabul, but provided onward movement of passengers from the region, while military aircraft were responsible for the evacuations out of Kabul. This greatly decreased the element of risk for the activated CRAF members. Additionally, the airlines faced no disruption to their commercial operations and qualified their participation in the program as a success, each putting out press releases on their involvement. CRAF activated without issue, and the carriers completed their task with no disruption to commercial business, and there was minimal risk for the participating crew (DoD, 2021). A potential challenge of the existing activation model is the lack of real-world experience beyond stage one activation. CRAF has never been mobilized to stage two or three, resulting in limited insight into how the system would perform under higher levels of demand. This gap creates uncertainty around both industry responsiveness and the government’s ability to carry out activation. Despite this, CRAF’s well-defined activation plans, clearly designated decision authority, and history of three successful stage one activations mean that the program would likely find success in any level of activation. **Assessment: High**

## National Armaments Consortium

### Overview

The National Armaments Consortium (NAC) is the largest cooperative entity partnering with the DoD to advance armament technologies in support of U.S. national security (Advanced Technology International, n.d.). The NAC aims to create a technological advantage over the United States’s adversaries by fostering collaboration between industry professionals, academia, and government (National Armaments Consortium, n.d.). Made up of over one-thousand companies, the NAC leverages research, development, and the acquisition process to expand U.S. military munition capabilities (National Armaments Consortium, 2019).

### Authorities and Policies

The NAC operates under the National Cooperative Research and Production Act (NCRPA) of 1993 (DoJ, 2025). The NCRPA, enacted by the U.S. Congress, strives to “promote research and development, encourage innovation, stimulate trade, and make and appropriate modifications in the operation of the antitrust laws” (FTC, 2019). Under this Act, joint ventures and standard development organizations, such as the NAC, can participate in cooperative activities without violating antitrust laws (DoJ, n.d.). The NAC has around 1,229 companies as members (Chatham House Interview, 2025). It is a public-private partnership overseen by an Executive Committee consisting of nine members from academia, small businesses, and large



businesses while the Executive Director guides the organization's daily operations including outreach and development initiatives. Complimenting the Executive Director and Committee, a Consortium Management Firm provided by the government called Advanced Technology International (ATI) assists with management and administrative tasks to ensure the needs of the NAC are met in an efficient manner (National Armaments Consortium, n.d.).

Within the NAC, there are three working groups which target specific armament strategies. Companies looking to contribute to armament solutions can apply to be a NAC Member Organization. There are members of the NAC that have foreign corporate ties through their U.S. subsidiaries, such as BAE Systems, which is a U.K.-based company with U.S. subsidiaries. However, an NAC membership is limited to U.S. companies or U.S.-based affiliates of foreign companies as all members must meet certain security requirements such as DD2345 which is a Military Critical Technical Data Agreement used by the DoD (National Armaments Consortium, n.d.). **Assessment: High**

### **Government-Industry Collaboration**

The NAC fosters strong government-industry relations. As a consortium, there is a formal channel for the DoD to communicate their needs. The consortium model exemplifies the power of many stakeholders pooling resources and coordinating to deliver capabilities for the DoD. The NAC offers a single point of access for the government to engage in a broad network of industry and academic partners. DoD agencies participating in armament technology outline their technical needs, which NAC members then address with solution proposals. Additionally, a secure, members-only NAC website provides networking capabilities and detailed insight into DoD armament technology strategies and requirements. This allows NAC members to communicate with DoD customers to collaborate on technological challenges and innovative methods (National Armaments Consortium, n.d.).

One of the benefits of the NAC is that it can plan and host conferences for government-industry engagement without the restrictions faced by government organizations. The NAC can provide a forum for such engagement where its members can attend for free and government representatives can meet with many potential industry partners over a short period of time. The NAC also promotes government-industry collaboration through surveys of industry. They can survey their member base and collect valuable feedback for the government, for instance for the House Armed Services Committee (Chatham House Interview, 2025).

The NAC provides solutions for the DoD through three Other Transaction Agreements (OTAs). These OTAs issue solicitations to which NAC members can respond. The solicitations describe what technology is needed, and NAC members can submit an Enhanced White Paper explaining their approach (National Armaments Consortium, n.d.). These NAC-only solutions are a key incentive for industry engagement. Through the NCRPA, the NAC is granted antitrust immunity which allows members to engage in research and development as well as production. Therefore, members can focus on prototyping and experimentation, which can spark more innovative practices. Together, these features help shorten the time from identifying a need to delivering production-ready capabilities (Federal Trade Commission, 2019). **Assessment: High**

### **Resources and Contracting**

Funding for the NAC relies on member organizations paying a \$500 annual membership fee. When a member organization receives a Project Agreement from the DoD, the NAC is subject to an award assessment (FTC, 2019). This assessment does not exceed more than 1% of the obligation amount and is reviewed throughout the contract period to help maintain the NAC operating reserve. The NAC supports faster development and production by giving DoD agencies access to rapid contracting mechanisms. These processes significantly reduce administration constraints, allowing prototypes, technologies, and manufacturing improvements



to move from concept to production more quickly than through traditional acquisition processes (National Armaments Consortium, n.d.). The consortium model also speeds up progress by improving coordination between industry and government, giving members early insight into requirements and enabling iterative feedback during development (Defense Acquisition University, n.d.). OTAs play a large role in continuing NAC sustainment. The use of OTAs provides a streamlined process that allows the NAC to bring their findings and prototypes to the Federal Government without being subjected to many regulations that apply under the Federal Acquisition Regulations (FAR). Through these OTAs, NAC member organizations are able to participate in efficient acquisition pathways, leading to a streamlined scaling process (National Armaments Consortium, n.d.).

The Federal Government has also increased its focus on building up the industrial base. Modernization and reform efforts have been undertaken to enhance the nation's future strategic environment (DoD, 2024). This indicates that the DoD has seen increased funding efforts focused on strengthening our industrial base (Obis, 2024). The NAC has grown significantly over the years, with member organizations hoping to take advantage of DoD contracts (National Armaments Consortium, n.d.). For example, obligations for ordnance and missiles, have increased tremendously since 2022 (Federal Procurement Data System, n.d.). The NAC should continue to see an increase in DoD activity, allowing member organizations to place their focus on innovation and development, contributing to the future of national armaments (Gerton, 2025).

#### **Assessment: Medium**

#### **Activation Phases and Utilization**

The NAC operates under OTA frameworks, which include defined milestones for prototype development and technology maturation. However, the NAC is more strongly driven by DoD supply-and-demand. DoD program offices sponsoring NAC projects provide requirements, timelines, and performance metrics for consortium members. While the NAC itself does not set policy, its projects follow a DoD-defined process, ensuring there is a government plan and measurable objectives for technological development. Through solicitation announcements, Requests for White Papers (RWP), and Requests for Proposals (RFP), posted on the NAC members-only website, member organizations can see upcoming opportunities and potential production needs (National Armaments Consortium, n.d.). This provides early visibility for U.S. industry, signaling when to scale workforce, manufacturing, or supply chain capabilities to meet anticipated DoD requirements. The DoD establishes requirements and timelines for member organizations providing a clear framework for utilization. **Assessment: High**

### **Munitions Campus**

#### **Overview**

The OUSD A&S Munitions Campus was established as a pilot project on September 26, 2023, by the Office of the Deputy Assistant Secretary of Defense for Industrial Base Resilience (DoD IBR) and their Manufacturing Capability Expansion and Investment Prioritization (MCEIP) office (DoW, 2023). The campus is currently being built and managed by the American Center for Manufacturing & Innovation (ACMI). The goal of the Munitions Campus is to provide space for hypersonics, energetics, advanced fabrication processes, and other defense industrial needs (American Center for Manufacturing and Innovation, n.d.). The Campus is located within the planned National Security Industrial Hub (NSIH) in Indiana, next to the Crane Army Ammunition Activity and Naval Surface Warfare Center (Kadzielski, 2025).

#### **Authorities and Policies**

The Munitions Campus is an example of a public-private partnership where part of the funding for the campus comes from DPA Title III, and the rest is provided by private capital



(Chatham House Interview, 2025). DPA Title III Section 303 allows the President to “create, maintain, protect, expand, or restore domestic industrial base capabilities essential for the national defense, the President may make provision” by several methods including “purchases of or commitments to purchase and industrial resource or a critical technology item, for Government use or resale” and “the development of production capabilities,” (U.S. Congress, 2014). The Munitions Campus is also part of MCEIP’s Pathfinders, which strives to maximize the DPA and Industrial Base Fund to quickly build the U.S. domestic industrial base and address critical supply chain challenges. Pathfinders strategies include using private capital, emphasizing market pull, creating “innovation clusters,” and expanding domestic production capability and capacity. The Munitions Campus incorporates each of the four strategies (OASD, 2025). MCEIP and DoD IBR are within the DoD Office of Industrial Base Policy, which carries out programs through its authorization from sections 2502 and 2505 of the U.S.C. Title 10 (DoW, 2023). Through their contract with the federal government, ACMI is building the infrastructure and facilities for the Munition Campus while also coordinating with future tenant companies and will be the manager for the site (American Center for Manufacturing and Innovation, n.d.). **Assessment: Medium**

### **Government-Industry Collaboration**

The Munitions Campus is a highly collaborative initiative between the federal government, the defense industry, private capital, and academia (Kadzielski, 2025). The Campus aims to address the problem that startup companies face in securing funding for real estate and facility construction. ACMI and their future tenant companies are using government funding augmented by private capital to secure land and construct facilities (Chatham House Interview, 2025). The Campus represents a method where the emerging defense industry, federal government, and capital market can work in tandem to increase access and eventually incorporate advanced manufacturing for munitions production.

ACMI currently has agreements with at least 16 companies including the solid rocket motor supplier Prometheus Energetics, considered the “anchor” tenant (Chatham House Interview, 2025; Kadzielski, 2025). Prometheus plans to use 550 acres of the NSIH for their headquarters and a production facility (Kadzielski, 2025). Kratos Defense & Security Solutions and RAFAEL Advanced Defense Systems, an Israeli company, established Prometheus to produce energetics, SRM, and other munitions while being based in the United States. Through the Munitions Campus, government and industry collaborate both in coordination and funding but also in location. The NSIH site (which includes the planned Munitions Campus) and the government’s Crane site are located directly next to each other. ACMI Federal and Crane Army Ammunition Activity signed an MOU in 2024 to integrate government agencies, researchers, and companies through a framework focused on the munitions supply chain, resilience of the domestic munitions stockpile, and development of new munitions technology (ACMI, n.d.). The Munitions Campus broke ground in February of 2026 (DoW, 2026).

The Munitions Campus is a promising example of government-industry collaboration and the larger NSIH site has surpassed projections for generating private capital funding showing how this model of collaboration could be effective. There is an anchor tenant company in place and agreements with a range of smaller companies. The proximity of the Crane and NSIH sites should also be instrumental in ensuring the success of Munitions Campus. However, it is too early to judge the effectiveness of the site before it is operational and the day-to-day functioning of the campus can be observed. **Assessment: Medium**

### **Resources and Contracting**

The Munitions Campus is partially funded by the federal government, through DPA Title III, and primarily funded through private capital. From the government, ACMI Federal won a





contract for \$50 million in funding for shared equipment and \$25 million in R&D funding for the Munitions Campus Pilot Program after undergoing a competitive process (DoW, 2023). ACMI was then able to start planning their NSIH site which is expected to generate \$600 million in private capital (Kadzielski, 2025). Initially, this endeavor was expected to generate half that amount, only \$300 million in funding, showing how the program has been able to surpass expectations (Chatham House Interview, 2025).

The Munitions Campus eases the burden of commercial companies by reducing lease expenses until companies begin pulling in revenue. Operational efficiency is also maximized while tenant resources can be put into production instead of construction. Available excess tenant space provides an easy transition to expand production and scale (ACMI, n.d.). This approach supports the scaling of manufacturing processes. Venture capital (VC) funding provides most of the funding for the Munitions Campus, but it could also potentially fund tenants on the campus who require specialized equipment. ACMI has signaled that VC firms are eager to provide additional funding to tenants for manufacturing equipment. VC firms and in-house finance companies find these investments to be particularly attractive (Chatham House Interview, 2025). The success of investment from private capital—far beyond initial expectations bodes well for the development of the Munitions Campus. So long as investments by private capital in the campus prove successful, funding should be relatively secure. **Assessment: Medium**

### **Activation Phases and Utilization**

The Munitions Campus is in the early stages of development with one anchor tenant secured and a range of smaller companies intending to be tenants. The system is in place to combine government funding with private capital and use private industry as the manager of the enterprise (Chatham House Interview, 2025). The Munitions Campus, as part of the Pathfinders initiative, could be considered “activated” now that the anchor tenant, Prometheus, has been confirmed, but the campus is too early in development for the program to be considered executed or yet a success.

The nascent program is also too early in development for any clear plan as to how tenants on the campus will be able to scale production quickly and effectively or how government will be able to direct such production increases during a conflict. There are initiatives in place that could support such activation, including a focus on using a shared location between the Munitions Campus and larger NSIH site and the government’s Crane campus to streamline production and supply chain resourcing. There is also a focus on integrating equipment and processes into campus facilities that use advanced manufacturing techniques. There is not a clear activation plan in the same way as a voluntary agreement, but there are early ideas around how the campus could grow or be replicated around other capabilities. ACMI intends to find more anchors and look to industries like drones, batteries, and pharma. The campus has a significant amount of space still left to fill, and it is expected that it would be refreshed over time with new tooling equipment. The design of the program, focus on innovation, and chosen location will likely contribute to the ability of Munitions Campus tenants to scale production, but concrete plans, a coordination framework for managing competing tenant needs, and on-the-ground equipment and facilities are still undetermined (Chatham House Interview, 2025). **Assessment: Medium**

## **Civil Reserve Manufacturing Network Case Study**

### **Overview**

The Civil Reserve Manufacturing Network (CRMN) is a public-private initiative recently established under the National Defense Authorization Act (NDAA) for FY2026 to ensure the





United States maintains a resilient industrial base capable of supporting defense requirements in times of national emergencies and conflict. CRMN leverages commercial manufacturing capabilities aligned with national security needs, ensuring rapid scaling of production for defense materials, components and systems (Albon, 2025). This initiative creates opportunities to formalize agreements and surge production in ways that go beyond traditional voluntary agreements (Chatham House Interview, 2026).

### **Authorities and Policies**

CRMN is authorized under section 1841 of NDAA FY2026, incorporating House bill (sec. 879) and Senate Bill (Sec. 220), which direct the Secretary of Defense to establish a working group under the Defense Industrial Resilience Consortium to support CRMN creation (Chatham House Interview, 2026). A CRMN Establishment Plan must be submitted to congressional defense committees within 120 days of enactment, and the first report is expected in June. The Secretary of Defense will also need to provide Interim Progress Reports on CRMN development for Congressional Oversight, allowing transparency for both government and industry stakeholders (U.S. House Rules Committee, 2025).

Congress has intentionally left significant decision-making authority to the Office of Industrial Base Policy (IBP), giving the DoD flexibility to determine how CRMN will operate in practice. As one official noted, “The proposal and how the activate has been the hardest part,” highlighting the ongoing process of figuring out the program’s operational details. This flexibility allows the CRMN to adapt its authority and activation approaches as lessons are learned during early implementation (Chatham House Interview, 2026). **Assessment: Medium**

### **Government-Industry Collaboration**

CRMN will establish formal, legally binding agreements between the DoD and private commercial manufacturers (Chatham House Interview, 2026). Industry participants can interact directly with the Pentagon, while the DoD provides surge requirements and technical standards to ensure readiness. By June 2027, the Department plans to enter into agreements with at least two manufacturers to participate in CRMN (Prusock et al., 2026). Implementation discussions highlight that CRMN is building on models like the CRAF but will move beyond voluntary participation. Companies can commit with legally binding agreements, creating on-call capacity and additional resources. The goal is for close cooperation between the organic industrial base and commercial industry. This is described as an “agile approach of iterations” with steep learning curves in the first few years, acknowledging the program will evolve as the department works through incentive structures, intellectual property considerations, and sector specific needs (Chatham House Interview, 2026).

CRMN participants will receive support for infrastructure and workforce development as well as recognition as critical contributors to national defense readiness. Additionally, participants may benefit from accelerated qualification and certification processes and, subject to appropriations, receive funding to speed up testing and modify existing manufacturing capabilities, with an emphasis on upgrading current facilities rather than building new ones. However, the incentive structure is still being developed and will expand as participation grows. Possibilities include investments in advanced manufacturing tools, models where the government buys capacity by placing tools in a facility, and other innovative mechanisms. CRMN planning has emphasized that the first few years will involve experimentation and iteration, noting the importance of aligning incentives with the diverse needs of different industries and sectors (Chatham House Interview, 2026). **Assessment: Medium**

### **Resources and Contracting**

CRMN has over \$177 million in funding from the NDAA to establish a network of commercial manufacturers capable of rapidly supporting defense production requirements.



Additional funds are allocated for Research, Development, Test and Evaluation (RDT&E) for the Army, Air Force, and other Defense agencies meant to expedite transformation when activated. Quarterly updates on the progress of establishing CRMN will be provided to congressional defense committees by the Assistant Secretary of Defense for IBP through September 30, 2027 (DoD, 2026). Interview insights highlighted that the program is not tied to any single contracting pathway, and the DoD can use a range of methods to fund and incentivize participants (Chatham House Interview, 2026). Coordination with other offices and attention to advanced manufacturing barriers, testing, and qualification processes are integral to CRMN planning (Prusock et al., 2026). **Assessment: High**

### **Activation Phases and Utilization**

CRMN is intended to maintain a pool of commercial manufacturers that can be activated to support defense production during wartime or national security emergencies. During activation, selected manufacturers would convert their production lines and increase output to provide surge capacity for critical defense needs under DoD direction (King & Spalding, 2026). CRMN will need a phased activation plan and clear triggers for production because it is expected that the program will move quickly once signals are sent. CRMN will also initially focus on domestic manufacturers, postponing allies or foreign companies to future phases. Component-level production will be easier to integrate into resilience models, and non-traditional or emerging manufacturers will be engaged progressively. The program aims to revisit its structure every few years to reflect lessons learned and evolving industry capacity (Chatham House Interview, 2026). **Assessment: Medium**

## **Findings: Elements of an Effective Manufacturing Security Program**

### **Authorities and Policies**

The case studies and supporting research interviews highlight the authorities and policies that best enable programs to succeed. Authorities like DPA Title VII Voluntary Agreements can be helpful but without regular use and experienced government staff and lawyers accustomed to creating them, they can be challenging to establish. The long gap that existed between voluntary agreements made it more difficult to create a new one when the Covid Voluntary Agreement was needed. Authorities need to allow for clear government-industry collaboration but also need to be used regularly so that the institutional knowledge and structures needed to create and maintain them are in place. Defense Production Act Title I requires the same regular use in order to be useful. Authorities and policies are most effective when they are in place well before a conflict. Anti-trust exemptions can be helpful for coordination between government and industry, but the legal steps to enact these exemptions take time. They should be used only when necessary to facilitate effective information sharing.

Organizational structures vary across case studies. The DOT's MARAD runs both VISA/MSP and VTA/TSP, the DOT and DoD share responsibilities for overseeing CRAF, and while the Voluntary Agreement was active, FEMA oversaw the COVID Voluntary Agreement and subsequent coordination efforts. Additionally, the DoD is responsible for CRMN. These five cases represent avenues for government-run programs, but there are public-private partnerships with different organization structures. For instance, a volunteer board is responsible for overseeing the NAC, while a private company, ACMI, established the Munitions Campus and will work with another company who will oversee daily operations. Organizational structures must be well articulated and involve only stakeholders and industry members relevant to the success of the program.



## **Government-Industry Collaboration**

Government and industry can coordinate in several productive ways as shown in the case studies. When information sharing is a challenge, authorities that allow for an anti-trust exemptions (DPA Title VII) can be helpful. These exemptions takes time to put in place, but when necessary, it allows government and industry to talk and share their respective needs and constraints. Anti-trust exemptions enable information sharing to an extent that is beneficial for collaborative activities focused on specific programs or industry sub-segments. For public and private entities to effectively coordinate, information sharing is key. In addition to authorities that enable communication, standing meetings and coordination timelines are also key. For the NAC, one method of coordination with the government included hosting conferences with industry for a government sponsor. For the Munitions Campus, coordination occurred between many different entities (the government sponsor, the coordinating company [ACMI], the tenant vendor [Prometheus], the potential smaller vendors) and the venture capital companies taking point on funding the endeavor. A regular pace of meetings and collaborations can lead to better trust and understanding between government and industry. The important caveat is that industry still requires clear incentives to participate. Programs must have a clear business case of increased peacetime contracts, regular payments for participation, or both.

## **Resources and Contracting**

The case studies demonstrate multiple options for funding and resources in addition to a range of contracting approaches. In the case of funding, certain programs (CRAF) are more expensive than others but can have Congressional support and a long history of success that helps maintain their funding. Initiatives like the Munitions Campus demonstrate a way in which the federal government can take a smaller role while private capital is responsible for a larger share of the initial and ongoing investment. These investments show industry commitment. The Munitions Campus was particularly efficient in pooling private capital and reducing the government's burden in financing the initiative. In addition to the need for continual, reliable funding, programs are also more successful with contract mechanisms that enable surge. Effective and established contractual relationships between industry and government are key.

## **Activation Phases and Utilization**

More than half of the case studies have outlined activation plans. These include VISA, VTA, CRAF, and the Covid Voluntary Agreement. As a consortium, the NAC does not have an activation plan, but there could be stages to growing the consortium or utilizing the consortium during times of conflict. The Munitions Campus is incredibly new, and there is a general idea of how the campus could grow, or the model could be replicated but these plans are not concrete, and the campus has only just broken ground. CRMN is also too new to have a clear activation plan. There are several factors that improve the success of an activation plan. One is that the plan has been activated before and any potential difficulties have been addressed. This is obviously not possible if there hasn't been a reason to activate the plan yet or it is new. A possible solution is consistent test exercises. There should also be specific operational constraints or stages to a conflict that trigger each activation stage.

## **Conclusion**

The United States recognized a gap in munitions production during the War in Ukraine that has only persisted as a burden for strategic planners. As a result of the conflict in Iran and the need to deter threats from Russia and China, the munitions problem could worsen. This issue is multifaceted as the United States and allies are unable to surge production both for the quantity of munitions needed and the breadth of munitions required to mount an effective deterrent strategy. Efforts to increase surge production require coordination between industry



and the government to find ways to build excess manufacturing capacity, uncover and address gaps in the supply chain, and utilize advanced manufacturing capabilities.

## References

- Albon, C. (2025, June 12). *House proposes manufacturing network to boost wartime arms production*. <https://www.defensenews.com/pentagon/2025/06/12/house-proposes-manufacturing-network-to-boost-wartime-arms-production/>
- Advanced Technology International. (n.d.). *National Armaments Consortium*. <https://www.ati.org/services/nac/>
- ACMI Group. (2025). *Home-ACMI*. <https://acmigroup.com/>
- Baldanza, B. (2021, September 1). *How the Civil Reserve Air Fleet benefits the country and helps airlines*. Forbes. <https://www.forbes.com/sites/benbaldanza/2021/09/01/how-the-civil-reserve-air-fleet-benefits-the-country-and-helps-airlines/>
- Busler, B. (n.d.). *FY20 NDAA fuel tanker study unclassified executive summary report*. U.S. Transportation Command. <https://www.maritime.dot.gov/sites/marad.dot.gov/files/2023-04/%28U%29%20FY20%20NDAA%20Fuel%20Tanker%20Study%20-%20Executive%20Summary%20Report%20%2815Jun21%29.pdf>
- Civil Reserve Air Fleet. (2024, February 23). *U.S. Department of Transportation*. <https://www.transportation.gov/mission/administrations/intelligence-security-emergency-response/civil-reserve-airfleet-allocations>
- Defense Acquisition University. (n.d.). *OTA today takeaways*.
- Defense Production Act of 1950. (2018, August 13). [https://www.fema.gov/sites/default/files/2020-03/Defense\\_Production\\_Act\\_2018.pdf](https://www.fema.gov/sites/default/files/2020-03/Defense_Production_Act_2018.pdf)
- Defense Production Act voluntary agreement report to Congress*. (2021, August 10). [https://www.fema.gov/sites/default/files/documents/fema\\_DPA-voluntary-agreement-report\\_2020.pdf](https://www.fema.gov/sites/default/files/documents/fema_DPA-voluntary-agreement-report_2020.pdf)
- Delegating additional authority under the Defense Production Act with respect to health and medical resources to respond to COVID-19. (2020, April 1). <https://www.govinfo.gov/content/pkg/FR-2020-04-01/pdf/2020-06969.pdf>
- Department of Defense Appropriations Act, 2026. (2026). <https://www.congress.gov/119/meeting/house/118889/documents/BILLS-1192ih.pdf>
- Department of Homeland Security. (2020, December 8). *Plan of action for PPE distribution under DPA*.
- DoD. (2021, August 22). *CRAF activation for Afghanistan*.
- DoD. (2024). *DoD releases first-ever National Defense Industrial Strategy*. <https://www.defense.gov/News/Releases/Release/Article/3643326/dod-releases-first-ever-national-defense-industrial-strategy/>
- DoJ. (n.d.). *Filing a notification under the NCRPA*.
- DoJ. (2025). *Notice pursuant to the National Cooperative Research and Production Act of 1993—National Armaments Consortium*. <https://www.govinfo.gov/content/pkg/FR-2025-04-21/pdf/2025-06796.pdf>
- DoW. (2023, September 26). *DoD establishes munitions campus pilot*.



- DoW. (2026). *Department of War announces groundbreaking of new munitions campus in Indiana*. <https://www.war.gov/News/Releases/Release/Article/4411124/departments-of-war-announces-groundbreaking-of-new-munitions-campus-in-indiana/>
- Exec. Order No. 13603. (2012).
- Federal Procurement Data System. (n.d.). <https://sam.gov>
- Federal Trade Commission. (2019). *National Cooperative Research and Production Act of 1993*. <https://www.govinfo.gov/content/pkg/COMPS-229/pdf/COMPS-229.pdf>
- FEMA. (2021, November 3). *Meetings to implement pandemic response voluntary agreement*.
- GAO. (2021, December). *COVID-19: Agencies are taking steps to improve future use of Defense Production Act authorities*. <https://www.gao.gov/assets/gao-22-105380.pdf>
- GAO. (2025, June 12). *Defense Production Act: Information sharing needed to improve use of authorities*.
- GAO. (2025, June 12). *Defense Production Act: Use and challenges from fiscal years 2018 to 2024*. <https://www.gao.gov/assets/gao-25-108497.pdf>
- Gerton, T. (2025). *The National Armaments Consortium is gearing up for 2026*. <https://federalnewsnetwork.com/contracting/2025/11/the-national-armaments-consortium-is-gearing-up-for-2026-with-fresh-direction-and-a-focus-on-armaments-innovation/>
- Harshberger, T. (n.d.). *Applying a Civil Reserve Air Fleet (CRAF) model to industrial production*.
- Joint Explanatory Statement to Accompany the National Defense Authorization Act for Fiscal Year 2026*. (2025, December). U.S. House Rules Committee. <https://rules.house.gov/sites/evo-subsites/rules.house.gov/files/documents/fy26-ndaa-joint-explanatory-statement-final.pdf>
- Kadzielski, P. (2025, April 9). *ACMI announces 16 companies to operate at munitions campus*. <https://acmigroup.com/2025/04/09/acmi-announces-16-companies-munitions-campus-nsih/>
- Kadzielski, P. (2025, March 26). *JLL to lead leasing ACMI's industrial hub*. <https://acmigroup.com/2025/03/26/acmi-properties-jll-nsih/>
- King & Spalding. (2026, January 12). *FY 2026 NDAA: Domestic sourcing, artificial intelligence, cybersecurity, and acquisition reforms*. <https://www.jdsupra.com/legalnews/fy-2026-ndaa-domestic-sourcing-9484065>
- Kok, A. (2021, February 16). *Tabletop exercise refines voluntary intermodal sealift agreement activation process*. U.S. Transportation Command. <https://www.ustranscom.mil/cmd/panewsreader.cfm?ID=B249EFB0-BDC7-3C1A-6403671EA44CC7CB&yr=2021>
- Lewis, I., & Coulter, D. Y. (2000). The voluntary intermodal sealift agreement. *Transportation Journal*, 40(1), 26–34.
- Lyman, M. (2011). *Civil reserve air fleet*. Air Mobility Command.
- MARAD. (2025, July 16). *Maritime security program (MSP)*. <https://www.maritime.dot.gov/national-security/strategic-sealift/maritime-security-program-msp>





Maritime Administration. (2024, October 18). *Voluntary intermodal sealift agreement*. Federal Register.

Maritime Administrator. (n.d.). *MSP & VISA facts*.

Maritime Administrator. (2005, September 22). *Maritime security program*. Federal Register.

Maritime Administrator. (2022a, November 7). *Renewal of the voluntary tanker agreement program; revised form of the voluntary agreement*. Federal Register.

Maritime Administrator. (2022b, June 6). *Tanker Security Program interim final rule*. Federal Register.

Maritime Administrator. (2022c). *VISA and MSP*.

Maritime Administrator. (2023, May 26). *Solicitation of applications for MSP agreements*. Federal Register.

McGinn, J., & Letts, O. (n.d.). *Options for strengthening the use of Defense Production Act Title VII*.

National Armaments Consortium. (n.d.). <https://www.nacconsortium.org/>

National Armaments Consortium Member Agreement. (2019, December 6).

Neenan, A. (2023). *The Defense Production Act of 1950*. Congressional Research Service.

OASD(IBP). (2025, July 22). *DoD tactics for engaging private capital—MCEIP pathfinders brief for NDIA ETI*.

Obis, A. (2024). *DoD dispelling lingering myths about OTAs*.  
<https://federalnewsnetwork.com/contracting/2024/02/dod-dispelling-lingering-myths-about-otas/>

Peter G. Peterson Foundation. (n.d.). What's in the CARES Act?

Peters, H. M. (2021). *Afghanistan evacuation: The Civil Reserve Air Fleet*.

Prusock, M. P., Schaengold, M.J., McQuade, P. F., Ross, E. M., & Malone, J. N. (2026). *The substantial impact of the FY2026 NDAA on federal procurement law*. Greenberg Traurig.  
<https://www.gtlaw.com/en/insights/2026/2/fy2026ndaa/fy-2026-ndaa-the-substantial-impact-of-the-fiscal-year-2026-national-defense-authorization-act-on-federal-procurement-law>

Tirpak, J. (2024, September 6). Still delivering after all these years. *Air & Space Forces Magazine*.

USA Spending. (n.d.). *Civil Reserve Air Fleet contracts FY2010–FY2025*.

Warren, M., & Conklin, T. (2024, September 4). *CRAF program benefits*.



# **The Gears of War: The State of the Aerospace Gear Industry and the Implications for Accelerating Acquisition**

Lieutenant Colonel Nathaniel L. Ross, USMC— CH-53K Integrated Product Team Lead, PMA-261.

## **Abstract**

Innovation is transformed into combat power through industrial capacity and production. Though much attention in Acquisition Reform centers on lowering barriers to entry for new entities and encouraging innovation and accountability for defense primes, these same primes are often paced by their sub-tier suppliers. This paper will explore how a subset of this supply base, namely aerospace gear and gear component manufacturers, illustrates the broader challenges of the defense acquisition enterprise. Chronic underinvestment from defense primes and national policymakers, coupled with recent inflationary trends, has created a supply base in this segment with flat or declining output and upward cost pressure. This paper will demonstrate, through exploration of United States Government Databases, that a buildup in the combat power required to defeat U.S. adversaries will be rate limited unless policy changes are made. Though the U.S. Government has tools to shore up supply chain capability, such as utilizing the Defense Production Act, these tools are often underutilized and/or too narrowly applied. Though the idiosyncrasies of the aerospace gear industry do not precisely generalize to the acquisition system as a whole, taking a finer, limited view of one aspect of the supply chain illustrates broader challenges and highlights the need to focused just as much on producibility as a catalyst for accelerating acquisitions.

**Keywords:** defense industrial base, defense economics, supply chain management, aviation manufacturing, defense production act

## **Introduction**

A hovering helicopter externally lifting a tactical vehicle, a tiltrotor transitioning to forward flight, a C-130 executing a short field landing. All these impressive and innovative feats of engineering and manufacturing contribute to the United States' primacy in the military domain. The precision raid in Caracas conducted by U.S. Special Operations Forces (Gordon, 2025), using rotorcraft from the 160th Special Operations Aviation Regiment, portray an image of inevitability and near invincibility; the equipment and operators of the U.S. Armed Forces will simply be there, and be ready, to accomplish any mission asked of them.

The reality of course is much more complicated, and messy. Setting aside for a moment the amount of training and rehearsals needed to conduct such an operation, think for a moment about what it takes to get an MH-47 off the ground. At the flick of a switch, digital computers provide signals to start fuel flow, power hydraulics, and activate avionics. Drilling in further, fuel ignites in the engine's combustion chamber, which powers a turbine. This turbine, spinning tens of thousands of revolutions per minute, connects to out output shaft which transfers the mechanical energy into a gearbox, where the real work begins.

The purpose of a gearbox sounds simple, but successfully achieving this purpose is anything but. At a rudimentary level, a gearbox reduces the output from the engine and changes the direction of this energy. By changing the ratio between speed and torque, propulsive power is created via gearbox output shafts which turn the rotor system, enabling flight. Within the gearbox itself, smaller gears, expertly machined, perform a tightly coupled dance to reduce engine speed by a factor of ten or more (Air Service Training [Engineering] Limited, 2007), under conditions where minor imperfections in the quality or craftsmanship performed on the metals could result in catastrophic failure, resulting in a loss of aircraft and aircrew. In many instances, an aircraft can safely land should one or more engines failure; however, a significant failure or seizure within a gearbox is more often than not unrecoverable.



This is not a paper about mechanical or aerospace engineering, but rather about defense acquisitions and the aerospace supply base. For years, commentators, academics, and policymakers have warned the country about fragility in the defense industrial base, to include landmark government reports such as the 2018 report “Assessing and Strengthening the Manufacturing and Defense Industrial Base and Supply Chain Resiliency of the United States” (Interagency Task Force in Fulfillment of Executive Order 13806, 2018). Indeed, as long as policymakers have recognized and publicized the China’s rise as a strategic competitor, there’s been a recognition that something is amiss in U.S. Defense Industrial Performance.

Two of the treatments recently put forward for what ails industrial performance are innovation and accountability. The war in Ukraine illustrates the vital role that innovation, both in tactics and in system development, has on battlefield success (Slusher, 2025). The argument for innovation as the path forward out of our industrial morass plays upon well-known U.S. industrial and technological strengths and is indeed the main argument in Seth Jones’ influential book *The American Edge* (Jones, 2026). As for accountability, recent Executive Orders and pronouncements by senior leaders place blame for poor industrial performance on misaligned incentives (Trump, 2026), in that defense primes plow profits back into share buybacks rather than investing in increased capacity or modernization.

While both treatments are necessary, their implicit focus on defense primes (for accountability) and new entrants (for innovation), leaves untouched the driving, propulsive force behind industrial output, namely the sub-tier suppliers. Without an understanding of and focus on the numerous often privately owned companies, that cast, forge, and machine the components comprising complex end products, policy makers risk sub-optimal or counterproductive outcomes. Much like the MH-47 uses the gearboxes in its transmission system to transform energy into propulsive flight, it is these sub-tier suppliers that transform the intellectual and fiscal resources of our country into combat power.

This paper will examine the economic background of a small subset of the supply chain, namely the sectors that comprise gearbox manufacturing. While their performance does not precisely generalize to the defense industrial base as a whole, important trends, and more importantly, recommendations, can be gleaned from relevant economic data. The picture this analysis reveals is one of fragility, where continuing business as usual is creating an unacceptable risk to success in future conflicts, where chronic underinvestment in Tier 2-4 suppliers has left key segments of the defense supply chain unable to rapidly productionize defense items in sufficient quantities needed to compete.

## Method

Using suppliers for CH-53K Gearboxes as a base, the six-digit North American Industry Classification System (NAICS) codes for these suppliers were cross-checked with the SAM.gov Entity database to ensure the NAICS codes were relevant to conduct broader analysis. See Table 1.



**Table 1. NAICS Codes Used in CH-53K Gearboxes**

| <b><u>NAICS</u></b> | <b><u>Description</u></b>                                      |
|---------------------|----------------------------------------------------------------|
| 331512              | Ferrous metal foundries noninvestment                          |
| 331529              | Other nonferrous metal foundries                               |
| 332111              | Iron forging                                                   |
| 332112              | Nonferrous forging                                             |
| 332510              | Hardware manufacturing                                         |
| 332710              | Machine shops                                                  |
| 332721              | Precision turned product manufacturing                         |
| 332999              | All other miscellaneous fabricated metal product manufacturing |
| 336411              | Aircraft manufacturing                                         |
| 336412              | Aircraft engine and engine parts manufacturing                 |
| 336413              | Other aircraft parts and auxiliary equipment manufacturing     |

## **Data Sources**

Using the above NAICS codes as a foundation for analysis, this study drew on five publicly available datasets across four U.S. government agencies to create a view of industry performance from 2012 until the present day. Though research focused on Tier 2-3 suppliers, represented by ferrous and nonferrous foundries (NAICS 331512, 331529), forgings (NAICS 332111, 332112), hardware manufacturers (NAICS 332510), machine shops (NAICS 332710), precision turned products manufacturers (NAICS 332721), and miscellaneous fabricated metal products manufacturers (NAICS 332999), Tier 1 aircraft manufacturing data (NAICS 336411–336413) were also analyzed. I chose 2012 as the base year, and the year to which all data was normalized, as the Defense Strategic Guidance (Sustaining U.S. Global Leadership, 2012) published early that year is one of the earliest strategic documents which speaks to an orientation away from counter-terrorism and towards the Pacific Competition.

Datasets used in this study included the Federal Reserve Bank of St. Louis' Federal Reserve Economic Data (FRED) Industrial Production and Capacity Utilization Indices via the FRED API. Six series were retrieved covering the period January 2012 through January 2026: aerospace and miscellaneous transportation equipment industrial production (IPG3364S), aerospace capacity utilization (CAPUTLG3364T9S), and industrial production indices for primary metals (IPG331S), fabricated metal products (IPG332S), machinery (IPG333S), and the producer price index for primary metals (PCU331331). Series were normalized to a January 2012 baseline (index = 100).

Bureau of Labor Statistics Quarterly Census of Employment and Wages (BLS QCEW). Establishment, employment, and wage data were retrieved for 2014–2023 using the BLS QCEW Open Data API. Annual data were extracted at the national level (area FIPS = US000) for private-sector employers (ownership code 5) across 11 six-digit NAICS industries. Variables included annual average establishment count, annual average employment level, average weekly wage, and total annual wages Bureau of Labor Statistics Industry Productivity Program (via FRED). Labor productivity (output per hour), unit labor costs, and capital input indices were retrieved from FRED, which hosts BLS Industry Productivity and Costs data (IPUEN series). Series were obtained for NAICS 331, 332, 333, 336, and 3364 for the period 2012–2024. All series are indexed to 2017 = 100 by BLS and were reindexed to 2012 = 100 for cross-industry comparability.



U.S. Census Bureau Annual Survey of Manufactures (ASM) capital expenditure and shipment data were obtained from two Census Bureau API endpoints. The timeseries/asm/industry endpoint provided data for 2012–2016, including total capital expenditures (CEXTOT), machinery capital expenditures (CEXMCH), building capital expenditures (CEXBLD), total receipts (RCPTOT), value added (VALADD), and employment (EMP). The timeseries/asm/benchmark2022 endpoint provided receipts, value added, and employment for 2018–2021. No ASM data were available for 2017 due to the Census Bureau's survey cycle, resulting in a gap in the capital expenditure series. Nominal values were deflated to 2012 constant dollars using the GDP Implicit Price Deflator (GDPDEF) retrieved from FRED.

Bureau of Economic Analysis GDP by Industry (BEA) nominal and real value added and gross output by industry were obtained from the BEA GDPbyIndustry dataset via the BEA API for the period 2012–2024. Four table IDs were used: nominal value added (TableID 1), real value added in chained 2017 dollars (TableID 10), nominal gross output (TableID 15), and chain-type quantity indexes for gross output (TableID 16). BEA industry resolution does not permit separate identification of NAICS 3364; aerospace data are reported as a combined “3364OT” (Aerospace and Other Transportation Equipment) category. A value-added share ratio (value added / gross output) was derived as an indicator of supply chain depth.

Granularity on the gearbox value stream varied from database to database. While it was possible to map specific gearbox suppliers to six-digit NAICS codes, in many cases data were only available at the 3-digit NAICS code level, which covered a much broader and potentially superfluous number of additional entities. However, and as stated previously, capturing a broader sense of the relevant industrial landscape met the needs of the study, in that it was assumed specific entities in the value stream were subject to the same trends and conditions as the broader, more expansive sector definitions. Further, more granular investigation may be warranted in future studies.

## Results

First, an examination of the FRED Industrial Production (Figure 1) and Industrial Capacity (Figure 2) data reveals a slight downward trajectory in most, but not all industrial categories. For the purposes of this study, the gearbox value chain to FRED Industrial category can be mapped as follows: Primary Metal Manufacturing encapsulated NAICS Codes starting with 331, meaning the mills and foundries (both ferrous and non-ferrous), Fabricated Metals consists of NAICS codes 332, comprising forging and machine shops, while Machinery Manufacturing consist of finished machinery as a product class NAICS codes 333. Finally, Aerospace consists of NAICS codes 3364 and can be mapped to Tier 2 and higher entities. While the most noticeable element of these datasets is the severe drop in both production and capacity due to COVID-19 shocks, an additional conclusion is that despite a decade or more of the re-industrialization drive, relevant segments of the gearbox value stream are producing less than they were when the re-balance to the Pacific began. However, all aspects of the value stream appear to be operated at 80% capacity or lower, indicating that capacity is not necessarily the constraint towards increasing throughput.








**Figure 1. Industrial Production Indices for Gearbox Value Chain Sectors, January 2012–January 2026 (Normalized to 2012 = 100). Data from Federal Reserve Bank of St. Louis FRED Database.**




**Figure 2. Industrial Capacity Utilization for Gearbox Value Chain Sectors, January 2012–January 2026. Data from Federal Reserve Bank of St. Louis FRED Database.**

Next, and perhaps more dire, we see the divergency in labor productivity and unit labor costs. As shown in Figure 3, in all instances there is a marked divergence in unit labor cost and productivity. As review, while Figure 1 indicates the total industrial output of the relevant sectors, productivity is a measure of output relative to inputs. Here, across the board, we see that it's becoming more expensive overtime to produce a comparable unit of output. Additionally, unit labor costs are increasing rapidly, which tracks the overall decline in productivity. While this topic will be discussed further later in this paper, the data shown here is likely no surprise to those working in defense acquisition, dealing with both schedule slips and cost overruns. Figure 4 combines both elements into one chart, clearly indicated the inflationary pressure across all sectors of the value stream.

Productivity vs. Unit Labor Costs by Industry (2012 = 100)




**Figure 3. Labor Productivity and Unit Labor Cost Indices for Gearbox Value Chain Sectors, 2012–2024 (Normalized to 2012 = 100). Data from Bureau of Labor Statistics Industry Productivity Program via FRED.**




**Figure 4. Combined Inflationary Pressure Indicators Across All Gearbox Value Chain Sectors, 2012–2024. Data from Bureau of Labor Statistics Industry Productivity Program via FRED.**

Next, an examination of the demographics of the value stream is in order. For this, an analysis of Bureau of Labor Statistics Data, specifically the number of establishments and the total amount of workers in each respective industry. Here an interesting split is seen, in that there has actually been an increase in the number of industrial establishments participating in the value stream, in some instances, increases of 27.1% (aircraft engines and parts), 32.7% (misc. fabricated metal products), and even 62.2% (aircraft manufacturing). This would seem to indicate significant reindustrialization. However, employment trends are more mixed. While certain industries do show an increase in employment, on the whole, total employment is down across the majority of industries. This additionally drives down the average number of employees per industry nearly across all sectors, further tempering optimism regarding reindustrialization. While this analysis focused on percentage changes for ease of comparison, a discussion of absolute values is warranted. In the case of non-ferrous foundries, those entities

which often produce the critical gearbox housings, we see that the workforce has increased by over 7% from 2012 until 2023, the last observation in this data set. That is an increase of only 853 workers, and while that's still a step in the right direction, it does not portend a massive shift is U.S. labor force trends.




**Figure 5. Percentage Change in Establishment Count by Industry Sector, 2012–2023. Data from Bureau of Labor Statistics Quarterly Census of Employment and Wages.**





**Figure 6. Percentage Change in Total Employment by Industry sector, 2012–2023. Data from Bureau of Labor Statistics Quarterly Census of Employment and Wages.**




**Figure 7. Average Number of Employees per Establishment by Industry Sector, 2012–2023. Data from Bureau of Labor Statistics Quarterly Census of Employment and Wages.**

Census data for Capital Expenditure Intensity, though extremely limited (2012 to 2016 only), is included here for completeness. Figures 8 and 9 show a flat and consistent Capex Intensity for the observation period, indicating a potential lag from strategic shift to Pacific competition towards the industrial actions needed to do so. We also see that the machinery share of CAPEX remained constant, with a notable dip in ferrous metal forging throughout the limited observation period.






**Figure 8. Capital Expenditure Intensity by Industry Sector, 2012–2016. Data from U.S. Census Bureau Annual Survey of Manufactures.**




**Figure 9. Machinery Share of Capital Expenditures by Industry Sector, 2012–2016. Data from U.S. Census Bureau Annual Survey of Manufactures.**

Last, we examine Bureau of Economic Analysis Data for Value Added and Real Output by Industry. Here, results are mixed. Though primary metals has seen its share of value added increase sharply from 2012, the output quantity has not seen a corresponding increase. Also notable is the decrease in both value added and output in fabricated metal and machinery.




**Figure 10. Value Added and Real Gross Output by Industry Sector, 2012–2024. Data from Bureau of Economic Analysis GDP by Industry dataset.**

## Discussion

What do each of these disparate data sets tell us about the state of the aviation gearbox industry, and perhaps for U.S. Defense production as a whole? In sum, the rhetoric about Great Power competition and reindustrialization has not been backed up with concrete changes. Instead, production appears flat or declining, investment is static, and costs are skyrocketing. The increase in labor costs was the most notable finding of this paper, and is also the most worrisome, as the labor force is on the whole decreasing across most of the sectors surveyed. It is important to recognize and distinguish that unit labor cost and labor salaries are not one and the same and instead refers to the compensation per hour worked divided by the output per hour, meaning unit labor cost could also be driven by falling productivity due to untrained workers, aging equipment, and rework. The result becomes a downward spiral for the Department of War, where rising costs decrease buying power, which tamps down demand, which further increases cost and results in exorbitant per-unit cost for aircraft. No wonder the shift towards lower-cost massed systems is so attractive now.

As a struggling sector with firm demand (U.S. Government contracts), skyrocketing cost, and poor output, the Tier 2-4 supply chain appears ripe for Private Equity's involvement, with their traditional focus on cutting cost and improving topline performance. Indeed, this has been the case over the last several years, with many smaller suppliers being purchased by Private Equity Firms, a trend exacerbated by the sector distress experienced during the COVID-19

Pandemic (Rowan et al., 2021). Though potentially a welcome development in the short-term, provided PE firms can execute on their strengths of driving efficiency and discipline into an otherwise struggling firm, one should question whether efficiency is the proper outcome in a reindustrialization push. Efficiency and cost control are often driven by reducing headcount and shedding unused, unneeded capacity, and we have already seen that many of the sectors surveyed are operating below capacity. However, an industrial base that is always and everywhere lean, efficient, and optimized for current demand is not one that can pivot to a mass build up, or really even produce systems in quantities necessary for credible deterrence.

Asking any company, let alone a publicly held company, to invest in unused capacity needed for a potential future contingency would, of course, be fiduciary malpractice. Shareholders expect that a company take actions that would maximize their return, and any CEO would not keep their job for long if they prioritized patriotic, yet uncertain investments at the expense of shareholder returns. So, while recent exhortations against companies prioritizing share buy-backs at the expense of contract performance is welcome (Trump, 2026), we must still realize that current incentive structures favor short term financial performance over long-term strategic national security considerations.

Further, while industry certainly has a huge role to play in national security, setting both the pace and direction of change is the responsibility of the U.S. Government. Tools such as the Defense Production Act, designed to “shape national defense preparedness programs and to take appropriate steps to maintain and enhance the domestic industrial base” (Levin, 2025) are available and indeed have seem to see somewhat of a renaissance following the Act’s extensive use in the COVID-19 Pandemic. However, while DPA Title III Authorities certainly could drive productivity gains through the purchase of industrial equipment, as well as more recently emphasizing the domestic production of critical minerals, labor force training and improvement programs are outside the purview of this act and instead fall under Department of Labor as well as a variety of state and local programs. This appears to be a legislative gap that could either be rectified due to an amendment to the DPA or else having an increased focus on interagency coordination between the Department of War, Department of Labor, and other relevant stakeholders.

This paper purposely began with a small, tailored scope, but as we have seen the implications and findings are anything but small and tailored. Practitioners in the Warfighting Acquisition system, particularly those operating at the program level, may lament the trends we see while also feeling a sense of helplessness. However, the recent advent of Portfolio Acquisition Executives (Office of the Undersecretary of War for Acquisition and Sustainment, 2025) provides an opportunity to affect change at a meaningful level. An individual program may focus on the gearboxes of their specific system, using a carrot or stick approach to an individual prime who is operating in an environment of misaligned incentives and negative demographic trends, and still suffer delayed schedules and cost overruns due to poor sub-tier performance. This is not a program-level problem, it is a structural problem. In PAE, with its broader scope and authority, there is a potential to make meaningful impacts across platforms and programs, to include consolidated and expanded DPA decision authority or advocacy that could shape and fix structural issues in their area of concern. As evidenced by increase senior leader attention, and increased defense funding, more and more energy is going towards correcting these structural issues. Without an appropriate transmission system, without the right gearing, tailored on discrete, measurable problems and outcomes, this energy will be misspent. Building an effective gearbox involves precision, and accelerating defense production requires a precise understanding of the challenges faced.



## References

- Air Service Training (Engineering) Limited. (2007). *Helicopter transmission systems*.
- Gordon, C. (2025, January 3). US airpower paved the way for Delta Force to capture Venezuela's Maduro. *Air & Space Forces Magazine*.
- Interagency Task Force in Fulfillment of Executive Order 13806. (2018). *Assessing and strengthening the manufacturing and defense industrial base and supply chain resiliency of the United States: Report to President Donald J. Trump*. Department of Defense.
- Jones, S. (2026). *The American edge: The military tech nexus and the sources of great power dominance*. Oxford University Press.
- Levin, A. G. (2025, June 12). *Evaluating the Defense Production Act. Statement before Financial Services Committee*. Congressional Research Service.
- Office of the Undersecretary of War for Acquisition and Sustainment. (2025). *Acquisition transformation strategy*.
- Rowan, L., Kumaria, R., & Campbell, B. (2021). *The private equity opportunity in aerospace & defense*. KPMG.
- Trump, D. (2026, January 7). *Prioritizing the warfighter in defense contracting. Executive Order*. The White House.



THIS PAGE INTENTIONALLY LEFT BLANK





## PANEL 8. WINNING THE COMPETITION: LEVERAGING ACQUISITION FOR STRATEGIC ADVANTAGE

Wednesday, May 6, 2026, 1445 – 1600 ET (1145 - 1300 PT)

### Panel Summary:

In an era of intensifying great power competition, defensive readiness alone is insufficient. This panel introduces the paradigm-shifting concept of "Offensive Acquisition"—leveraging procurement processes to actively target adversary weaknesses and strategic asymmetries. By exploring the behavioral drivers behind successful acquisition acceleration in Norway and identifying transferable domestic innovations for U.S. Foreign Military Sales (FMS), these researchers demonstrate how individual expertise and agile methodologies can be weaponized to maintain global military preeminence.

**Chair: Dr. Dave Lewis, VADM USN (ret.)**, Professor of the Practice, Naval Postgraduate School, former Director of the Defense Contract Management Agency (DCMA)

### Presenters:

**From Policy to Practice: Drivers of Acquisition Reform** – *Ane Ofstad Presterud, Norwegian Defence Research Establishment*

**Offensive Acquisition: How Acquisition Processes Can Seize the Initiative in the US-China Competition** – *Emily de La Bruyere, Horizon Advisory*

**Defense Acquisition and Contracting Approaches: Implications for Foreign Military Sales**  
*Jerry McGinn, Center for Strategic and International Studies (CSIS)*



**Dr. Dave Lewis, VADM USN (ret.)**—is a retired United States Navy Surface Warfare Officer and Engineering Duty Officer. He currently serves as a Professor of the Practice at the Naval Postgraduate School, Monterey, CA as an intermittent employee. He works full time in the Defense Industry.

His leadership assignments included duty as the Aegis Shipbuilding Program Manager, Program Executive Officer, Ships, Commander, Space and Naval Warfare Systems Command, and prior to his retirement, as Director of the Defense Contract Management Agency.

His afloat tours were in USS Spruance (DD-963), USS Biddle (CG-34) and USS Ticonderoga (CG-47), all in combat systems and communications assignments.

David Lewis earned bachelor's and master's degrees in computer science and a certificate of completion in Weapons Engineering. He recently completed his Ed.D in Leadership and Learning in Organizations. He is a graduate of the Naval War College, Vanderbilt Peabody College, the Naval Postgraduate School and the University of Nebraska.

He is a lifetime Member of the US Naval Institute, the American Society of Naval Engineers and the Surface Navy Association and has published several articles on shipbuilding and warship engineering topics.



# From Policy to Practice: Drivers of Acquisition Reform

**Ane Ofstad Presterud**—Research Scientist at The Norwegian Defence Research Establishment (FFI) and Head of FFI's INLOG research program. Visiting Research Scholar at NPS from August 2025 till July 2026. [ane-ofstad.presterud@ffi.no]

**Frida Waage Prebensen**—Research Scientist at The Norwegian Defence Research Establishment (FFI) and the INLOG Research Program. [Frida-waage.prebensen@ffi.no]

**Jan Erik Voldhaug**—Research Scientist at The Norwegian Defence Research Establishment (FFI) and the INLOG Research Program. [jan-erik.voldhaug@ffi.no]

## Abstract

The current geopolitical environment has intensified security concerns and increased pressure to deliver military capability rapidly. A persistent challenge has been the long timelines associated with defense acquisitions. In Norway, successive governments have sought to improve the effectiveness of defense acquisitions by promoting greater use of off the shelf solutions as the preferred acquisition strategy, rather than costly custom-developed solutions. Despite repeated policy signals, implementation remained limited for years. Recent evidence, however, indicates change, with a growing share of new acquisition projects adopting an off the shelf approach. This study examines what has driven this shift and what it can tell us about implementing acquisition policy more broadly. Drawing on policy implementation theory, we develop hypotheses about the conditions influencing the implementation of off the shelf policy. Empirically, we combine document analysis, previous research, and analysis of survey data. The findings align weakly with explanations based on changes in market conditions, governance structures, incentives, or policy communication. Instead, the observed change aligns more closely with shifts in the attitudes and perceptions of implementing actors. We argue that the policy has become more meaningful as its justification has shifted from a primary focus on cost efficiency toward enabling rapid build up of defense capability.

## Introduction

In an increasingly unpredictable geopolitical environment, the ability to rapidly procure, integrate, and adopt new defense materiel is critical. Yet, both Norwegian and international experience show that this remains difficult. Despite repeated reforms, acquisition processes are still lengthy and delays are widespread. Designing an effective defense acquisition system is challenging in a context where projects differ substantially and framework conditions continuously change. Ongoing adjustments to acquisition arrangements reflect the need for refinement but also illustrate how political directives do not always translate into practice. Ultimately, success depends not only on policy design, but on how policies are implemented.

For years, Norwegian authorities have promoted a policy of procuring more off-the-shelf solutions, arguing that this could accelerate defense acquisition and reduce costs. Implementation, however, progressed slowly. The Norwegian Armed Forces continued to rely heavily on costly custom-developed solutions, and acquisition projects remained long-running and often delayed.

Only in recent years has a measurable shift occurred. Norway is now buying more off-the-shelf solutions. Rather than assessing whether the current level of off-the-shelf procurement is the right one, we examine whether observed developments indicate a shift in acquisition decision-making in line with stated policy objectives. Drawing on policy implementation literature we analyze the actual implementation of the policy in the defense acquisition system over time, to understand the underlying conditions that shape implementation success. The goal is to generate knowledge that can support broader acquisition reform.



Although Norway is a small country with limited military capacity compared to major powers such as the United States, the findings have wider relevance. The study sheds light on general implementation challenges in complex acquisition systems faced by both large and small countries. The study may also contribute to the policy implementation literature by showing that implementation change can occur in the absence of structural reform, and by linking such change to shifts in how implementing actors perceive and justify the policy.

## Research questions

1. What conditions have shaped the implementation of the policy promoting increased use of off-the-shelf solutions within Norway's defense acquisition system?
2. What lessons can be drawn from the implementation of this policy for the design and execution of future reform initiatives?

## Background

A consistent finding in research on policy implementation is that achieving change is exceedingly difficult. It is particularly challenging to bring about change across multiple organizations or levels of government, a characteristic often seen in reform efforts within the defense sector (McLaughlin, 1987). In the Norwegian defense sector, the challenges associated with lengthy acquisition processes and delays have been highlighted for years (Prebensen & Voldhaug, 2026; Presterud & Øhrn, 2015; Riksrevisjonen, 2005, 2024). Several changes to both organizational structures and processes have been attempted, but the challenges remain unresolved. Similar difficulties are also reported in other countries. (GAO, 2024, 2025; House of Commons Defence Committee, 2023; National Audit Office, 2021).

In this study, we examine the implementation of a specific acquisition policy to identify the conditions that have shaped it. We use *policy implementation* to mean what occurs between "the establishment on the part of government to do something, and the ultimate impact in the world of action" (O'Toole, 2000). The policy examined is the Norwegian government's policy to use off-the-shelf solutions as the preferred acquisition strategy in acquisition projects for the Norwegian Armed Forces. In the following section, we provide a brief background for the policy and its implementation.

After the end of the Cold War, defense was deprioritized in Norway, as in many other countries. While defense budgets declined, the costs of defense equipment increased, making it increasingly difficult to modernize the Armed Forces at the required pace. In response, authorities introduced adjustments to the acquisition system. One key measure was to promote greater use of fully developed, off-the-shelf materiel. The aim was to enable reduced time and costs by limiting nationally tailored development and customization. Because acquisition projects are primarily designed to meet project-specific requirements, increasing the use of off-the-shelf solutions did not occur without intervention and had to be articulated as a distinct policy objective.

This policy did not constitute a major or disruptive reform for the sector. Nevertheless, it required changes to established practices, organizational culture, competence, and decision-making processes. It affected, and still affects, all key actors in the acquisition system. At the same time, because the policy was relatively concrete, delimited and measurable, it represents a useful case for studying how policies are implemented within the acquisition system.

In 2015, the Norwegian MoD asked The Norwegian Defense Research Establishment (FFI) to study the use of the policy within the Armed Forces (Presterud et al., 2016). In the



study, we developed a conceptual framework and a method for assessing the share of off-the-shelf acquisition projects in the acquisition portfolio. The assessment has since been repeated twice, in 2019 and in 2023 (Berg & Waage, 2021; Voldhaug et al., 2024). This enables us to trace the implementation of the policy over time and assess how different conditions relate to the observed changes. The results of the assessments are shown in Figure 1.




**Figure 20. Share of Off-the-Shelf Acquisition Projects in the Norwegian Defense Acquisition Portfolio in the Period 2015–2023 (Voldhaug et al., 2024).<sup>1</sup>**

The figure shows the share of projects in the defense acquisition portfolio classified as off-the-shelf. There were only minor changes between 2015 and 2019, followed by an increase from 2019 to 2023. While this increase is not particularly large, this is partly due to the long duration of acquisition projects, as many projects classified in 2015 and/or 2019 were still part of the portfolio in 2023. When considering only projects that entered the portfolio after 2019, as many as 54% are classified as off-the-shelf (Voldhaug et al., 2024). This indicates that a change has taken place in the implementation of the government's off-the-shelf policy after 2019.

A major challenge for evidence-based development of the acquisition system is that reforms are introduced and revised faster than it is possible to evaluate their effects. Our study is an attempt to examine how the Norwegian defense sector has implemented and responded to new policy over time. The findings from this study cannot be straightforwardly generalized to the implementation of larger and more disruptive reforms. Nevertheless, the analytical strength of the case lies in the fact that the policy is concrete, delimited, and has remained stable over time, allowing us to trace implementation dynamics and discuss possible causal mechanisms. Hopefully, the insights we provide can help inform future reform efforts.

## Delimitation

A successful policy is one that achieves its stated objectives. What constitutes successful policy implementation is a distinct analytical question. Following Winter (2002/2005) we distinguish between studies of policy outputs and policy outcomes. We use *outputs* to refer to the immediate products of implementation, in our case more off-the-shelf solutions in the acquisition project portfolio. *Outcomes*, on the other hand, refer to the subsequent effects of these outputs, such as whether acquisition would become faster, cheaper, or delivers the

<sup>1</sup> We measure the number of projects rather than their monetary value because our interest lies in the decisions made in each individual project. Had we instead measured the share in financial terms, it would have been considerably lower, since the largest acquisition projects in Norway are predominantly development projects.

intended capability. This study focuses on the outputs, and specifically on how acquisition policy is translated into decision-making practice.

The observable changes in the acquisition portfolio have been modest. This is not unexpected given the long duration of most defense acquisition projects and high path-dependence. In the context of this study, implementation success is therefore understood primarily in terms of direction in decision-making and underlying decision-making logic. Even limited changes in the composition of the portfolio are analytically meaningful as they reflect shifts in actions at key decision point.

## Method

In this study, we develop hypotheses about conditions that influence implementation of the off-the-shelf policy in acquisition projects, drawing on the literature on policy implementation theory. We then explore these hypotheses using empirical data based on document analysis, including previous research, and data from surveys.

### Document Analysis

To examine the structural conditions and their development over time, such as how the policy has been communicated, the processes through which it has been implemented, and the associated incentive structures, we conducted document analysis. By this we mean a structured review of the public documents that describe the government's policy and objectives for the period 2004–2025. All documentation analyzed is unclassified and publicly available, except for the documents on performance indicators derived from classified agency reporting.

In addition, we have drawn on information and research findings from previous FFI studies on defense acquisitions and efficiency reforms in the Norwegian defense sector.

### Analysis of Survey Data

In analyzing attitudes and perceptions related to the policy, we draw on data from two earlier surveys conducted among actors involved in the acquisition process in the Norwegian defense sector, one from 2014 and one from 2023. The surveys were originally carried out in connection with previous studies on related topics, but parts of the data have been analyzed and used in new ways in this study. Table 1 summarizes key information about the survey.<sup>2</sup>

**Table 5. Key Information About the Surveys in 2014 and 2023.**

| Year        | 2014                                                                                                                                    | 2023                                                                                    |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| Respondents | 119                                                                                                                                     | 98                                                                                      |
| Method      | Audience response survey (anonymized responses)                                                                                         | Online survey (anonymized responses)                                                    |
| Agencies    | Defense Staff, the service branches, The Armed Forces Logistics Organization (later the Norwegian Defense Materiel Agency) and the MoD. | Defense Staff, the service branches, the Norwegian Defense Materiel Agency and the MOD. |

<sup>2</sup> More details about the data and the survey methodology can be found in Presterud and Øhrn (2015) and Voldhaug et al. (2024).



## Methodological Limitations

The study has several methodological limitations. We have not had access to internal documents or communication between the actors, nor to classified information about the acquisition portfolio or the execution of the acquisition projects. This means that we are unable to capture directives that may have been given through individual taskings or internal dialogue.

The survey data also has certain weaknesses. As with all surveys, it is hard to know whether the respondents are representative of the entire population or whether the sample is biased. We also cannot know whether the respondents answered truthfully, even though they were informed that their responses would be anonymized. In addition, both surveys were conducted in connection with other studies, meaning that the questions were not tailored to the research questions of this study. This limits our ability to analyze the actors' attitudes and perceptions as thoroughly as desired. The sample is also too small to allow for robust statistical analysis.

These limitations prevent strict causal identification. Nevertheless, the data are unique in Norway as they show how actors' perspectives have evolved over ten years. Although the number of respondents is small, it is important to emphasize that the acquisition workforce in the Norwegian defense sector also is small. Those who responded therefore represent a significant share of those working with acquisition projects on a daily basis. However, the sample may still be skewed due to nonresponse and self-selection, and we cannot rule out systematic differences between respondents and non-respondents.

## Theoretical Background

"Policy success and effectiveness depend on the implementation process as much as it depends on the policy design that stands at the beginning of the implementation process."  
(Sager et al., 2024).

Policy implementation theory as a scholarly field emerged from this recognition in the 1970s. Many point to the classic study by Pressman and Wildavsky (1973) as the real starting point for the literature in this area. Here, the authors examined the implementation of a development project in California and highlighted why seemingly sound political decisions often fail during the implementation phase. Since then, the field has expanded considerably. The purpose of implementation studies is to provide insight into which conditions can be influenced in order to ensure that policy is actually put into practice (Van Meter & Van Horn, 1975). In this study, we draw on the literature to develop hypothesis of the implementation of acquisition policy in Norway.

Policy implementation models have been applied to a wide range of policy domains, including public administration, environmental and energy policy, health policy, and education policy (Sætren, 2024). Researchers in the field often focus on major political reforms and their effects on the population. Policy rollout within defense acquisition differs somewhat from this: the number of people involved is usually smaller, and there are fewer layers between the decision-makers and those responsible for implementation. The target group is also not an entire population, but the Armed Forces themselves. Also, in contrast to many other societal reforms, implementing acquisition reform is not necessarily very costly. This means that some of the themes frequently emphasized in policy implementation theory, such as financing of implementation, are not particularly relevant for our study.

The policy implementation literature has traditionally been divided between those who examine implementation from the perspective of decision-makers and those who analyze it from the viewpoint of the implementing actors. While researchers with top-down perspective aim to provide normative advice to decision-makers about the implementation process, scholars in the



bottom-up school focus more on describing the conditions that may hinder the achievement of policy goals (Matland, 1995). The aim of this study is to support decision-makers in the defense sector by improving understanding of the conditions that shape policy implementation. To this end, the analysis primarily builds on top-down perspectives in the policy implementation literature, which focus on how policy design and governing arrangements shape implementation outcomes.

As its starting point, our analysis takes the overarching framework of Sabatier and Mazmanian (1980) which we adapt to the context and policy examined here. We investigate three overarching explanatory conditions shaping implementation: (1) conditions related to the tractability of the problem, (2) structural conditions, and (3) non-structural conditions. The remainder of this chapter elaborates on these categories and presents the associated hypotheses, which are explored in the subsequent analysis.

### **Tractability**

An important condition for policy implementation is how challenging the problem to be addressed actually is. This category differs from the next two in that it concerns external conditions that influence implementation, conditions that may be difficult for the actors involved, whether decision-makers or implementers, to influence. Sabatier and Mazmanian (1980) refer to this category as “the tractability of the problem,” meaning the extent to which the problem addressed by the policy is in fact manageable. In our case, we limit this to the degree of difficulty and the time required to put the policy into practice, that is, to procure more off-the-shelf solutions.

We examine tractability primarily to assess whether changes in external conditions could plausibly account for the observed implementation shift, that is, the increased tendency after 2019 to select off-the-shelf solutions, rather than as a fully testable explanatory mechanism in its own right. Against this background, we examine the following hypothesis as a potential contextual explanation:

- The observed development in the implementation of the policy may be related to the problem having become more tractable over time.

### **Structural Conditions**

This category of conditions concerns the structural framework established by decision-makers for the policy, namely how the policy is communicated, the processes and governance structures surrounding it, and the incentive arrangements that accompany it. These conditions are typically under the control of decision-makers and are therefore often a central focus when new policy is to be implemented.

The top-down literature emphasizes that clear communication is essential for successful implementation. Goals and guidelines must be precise, consistent, and comprehensible for implementation to be effective (Sabatier & Mazmanian, 1980; Van Meter & Van Horn, 1975). Not only must the policy itself be clear, but the implementing actors must also know what is expected and required of them. Vague mandates and guidelines may lead to competing perspectives or dominant interest groups shaping the execution of the policy (McLaughlin, 1987). Matland (1995) highlights that ambiguity, both regarding objectives and the means for achieving them, can lead to misunderstandings and uncertainty, but he also notes that vague goals may reduce conflict when new policy is being designed.

Furthermore, the process for policy execution and the governance structure of the policy play a vital role. Long decision chains, numerous actors, and multiple veto points increase the likelihood of implementation failure (Pressman & Wildavsky, 1973; Sabatier & Mazmanian, 1979). Hupe (2011) refers to this as the Thesis of Incongruent Implementation. A policy process



usually involves several actors at different administrative levels. Hill and Hupe (2003) introduced the concept of the multi-layer problem to describe the challenge that arises when a policy must pass through several layers of governance before it can be implemented and executed. Different actors apply their own assessments and priorities in implementation, which may affect both output and outcome.

The literature also highlights the role of incentives in implementation. Situations in which actors are responsible for implementing a policy on behalf of a principal can be viewed as a principal–agent problem. Incentives are a tool for addressing this challenge by motivating actors to behave in a particular way. Successful implementation may depend on mechanisms that enable decision-makers to induce the behavior they desire. The greater the burden associated with implementing the new policy, the stronger the incentives typically need to be (McLaughlin, 1987; Van Meter & Van Horn, 1975).

In the empirical analysis that follows later, we assess whether strengthened structural conditions, namely communication, processes, governance structures, and incentives, have shaped the observed development in implementation of the off-the-shelf policy. Accordingly, we examine the following hypothesis:

- The observed development in implementation of the off-the-shelf policy in acquisition projects may be attributable to strengthened structural conditions.

### **Non-structural Conditions**

In the framework of Sabatier and Mazmanian (1980), non-structural conditions cover the non-legal factors that influence whether a policy is implemented. We limit our discussion to the conditions that concern the implementing actors. The structure established by decision-makers shapes how the policy is supposed to be carried out, but if the actors responsible for implementation oppose the policy, success is unlikely. Those who actually implement policy are individuals with their own interests and values that influence the policy in practice (Sager et al., 2024). The bottom-up literature rests largely on the assumption that the preferences and interests of the implementing actors must align with the policy for implementation to be successful (Sætren, 2024). The actors' sense of meaningfulness in relation to the policy influences their willingness to implement it, which in turn affects how successful implementation will be (Tummers et al., 2012). If actors perceive the policy as meaningful for society or the group it targets, this will strengthen their willingness to implement it.

The experience of the implementing actors also plays a role. Giacchino and Kakabadse (2003) emphasize that the actors' skills, defined as, among other things, experience, appropriate training, and the right qualifications, are key criteria for successful implementation. Employees who are overworked or who lack adequate competence or training, on the other hand, may hinder implementation (Van Meter & Van Horn, 1975).

We examine whether changes among the implementing actors align with the observed development in implementation of the off-the-shelf policy. Accordingly, we advance the following hypothesis:

- The observed development in implementation of the off-the-shelf policy in acquisition projects may be attributable to changes among the implementing actors.

### **Analysis**

Drawing on the literature, we developed three hypotheses that structure our analysis of the observed development in implementation of the off-the-shelf policy in acquisition projects over the past decade. We thus distinguish analytically between the policy implementation output and the conditions that may explain the development in this output. In this section, we use the



available empirical evidence to explore the hypotheses. The hypotheses are not tested as competing causal propositions in a strict sense, but as alternative explanatory perspectives on the observed patterns of the implementation of the policy over time. Our assessment therefore focuses on how well each perspective helps make sense of the empirical developments observed, rather than treating the hypotheses as symmetrically testable propositions. We begin by examining whether changes in external conditions (tractability) could plausibly account for the observed development, before turning to the other two categories of conditions more directly related to how the policy is governed and implemented.

## **Tractability**

Implementing minor policy changes aimed at addressing relatively simple problems is easier than implementing major policy changes intended to solve complex problems. Likewise, policies with long lead times will not produce effects until several years have passed. We have already established that the policy change examined in this study is relatively limited: actors are expected to select the off-the-shelf acquisition strategy more frequently. Nevertheless, the challenge of implementation should not be underestimated. Regardless of both structural and non-structural conditions, achieving the goal of increased use of off-the-shelf solutions is difficult if such solutions are not actually available or if acquiring them takes years. In this section, we examine whether the increase in off-the-shelf acquisitions observed between 2015 and 2023 may be due to the strategy having become easier to pursue. We assess both the actual availability of off-the-shelf solutions and the role of time.

A thorough analysis of whether off-the-shelf availability is greater now than ten years ago would require a broader market study of the defense industry in relevant countries. This has not been feasible within the scope of this study. We therefore limit the assessment to examining key high-level indicators of developments in the defense industry and the Norwegian Armed Forces' needs. Since 2015, global defense spending has increased every year, driven in part by Europe's intensified defense efforts following Russia's attack on Ukraine (Liang et al., 2025). This surge in demand has boosted revenues in the defense industry, but has also resulted in order backlogs and long lead times (Scarazzato et al., 2025).

How has this development affected Norway's access to off-the-shelf defense materiel? The substantial industry growth in recent years may have increased access to fully developed materiel, while strong demand may simultaneously have constrained it. On this basis, we cannot determine the net effect on Norway's access to off-the-shelf materiel.

We can, however, assess the demand side and examine whether the Norwegian acquisition portfolio has become more "off-the-shelf-friendly" over time. When comparing the overview of planned acquisitions in the Armed Forces in 2014 with the corresponding overview in 2023, there is little to suggest such a development (Forsvarsdepartementet, 2014, 2023a). ICT systems are often highlighted as an area with strong potential for off-the-shelf procurement, yet the share of ICT acquisitions was, perhaps somewhat surprisingly, just as large in 2014 as in 2023. The major difference between 2014 and 2023 lies in the maritime and air domains. Whereas half of the portfolio in 2014 consisted of air systems and only 14% of maritime systems, the two categories are now equal in size (30% each). Traditionally, the air domain has been an area in which Norway could procure off-the-shelf solutions from allies, while the maritime domain has been characterized by a need for tailored solutions (Presterud et al., 2016). Thus, there is little in the overall composition of the portfolio to suggest that using the off-the-shelf acquisition strategy is easier now than ten years ago.

Time is also an external factor affecting implementation. It takes time for political decisions to translate into real change. According to Sabatier and Mazmanian (1979) one must expect at least three to five years to ensure adequate compliance with a new policy. Given the



lengthy timelines of acquisitions in the Norwegian defense sector, it is reasonable to expect that the observable changes in implementation outputs of a new policy will only become visible after several years.

The policy examined in this study was introduced in the mid-2000s. From the first measurement in 2015 to the last in 2023, there has been a positive development in policy implementation, which may indicate that the system and the organization required time to move from decision to actual implementation. However, the development has not been uniformly positive throughout the period. From 2015 to 2019, the share of off-the-shelf acquisition projects declined somewhat, before increasing from 2019 to 2023. Taken together, the uneven development suggests that the overall trend is unlikely to be explained primarily by time alone.

Drawing on the research literature, we proposed the hypothesis that the tractability of the off-the-shelf policy, how difficult or time-consuming it has been to implement, could plausibly account for the observed developments in implementation. Based on the available data, we examined whether implementing the policy (i.e., selecting more off-the-shelf solutions at key decision points) has become easier over time in ways that might explain the observed shift. Overall, the evidence provides only weak support for the proposition that changes in tractability are associated with the development.

### **Structural Conditions**

Having assessed whether changes in external conditions (tractability) could plausibly account for the observed development, we now turn to conditions more directly related to how the policy is governed and enacted in practice. In this first section, we examine whether changes in the structural conditions surrounding the policy align with the observed development in its implementation. We begin by looking at communication, which the literature emphasizes should be clear and consistent.

It is not possible to identify exactly when increased use of off-the-shelf solutions as an acquisition strategy was first communicated as a political objective, as the policy has gradually evolved from vague formulations to more explicit ambitions over many years. It is nevertheless evident that successive governments, from across the political spectrum, have expressed a desire for greater use of off-the-shelf solutions in the Norwegian Armed Forces. To trace how policy communication has developed over time, we reviewed key governing documents for the Norwegian defense sector dating back to the early 2000s.

The government's long-term plans in the years from 2004 to 2012 contain references to the use of off-the-shelf solutions in defense acquisitions (Forsvarsdepartementet, 2004b, 2008b, 2012a). However, the policy received limited attention in these documents and was also scarcely mentioned in the documents that operationalize the strategy (Forsvarsdepartementet, 2004a, 2008a, 2012b). No specific targets or milestones were established. In other words, the policy does not appear to have been a priority during this period.

Around 2015, the government started expressing the policy more explicitly, with guidance stating that the acquisition of fully developed products should be the default approach (Forsvarsdepartementet, 2015, 2016a). The MoD also introduced a performance indicator related to the policy, albeit without accompanying targets or milestones (Forsvarsdepartementet, 2016b). Regardless of the changes in communication, the 2019 assessment showed a decline in the use of off-the-shelf solutions, suggesting that clearer communication did not result in increased implementation.

In the 2020 plan, the policy was mentioned less frequently than before, and the argument concerning cost savings was de-emphasized in favor of an emphasis on faster acquisition timelines, (Forsvarsdepartementet, 2020). The performance indicator disappeared,





and the Armed Forces remained without concrete objectives or milestones on the policy. In the years that followed, policymakers continued to promote off-the-shelf acquisition as a policy in this manner, sporadically and without explicit requirements or expectations. The emphasis on faster acquisition timelines persisted.

Overall, the policy has been articulated in a general and somewhat inconsistent manner, with varying rationales over time. Because our review is based on public documents, it may not capture all informal guidance or lower-level directives. Within these limits, the document review aligns only weakly with the proposition that clearer communication, with explicit goals and milestones, has been a key driver of the observed shift in the implementation of the off-the-shelf policy.

Further, we examine how the authorities have organized the policy's processes and governance structures, which the literature identifies as important for effective implementation. Policies of this type are decided by the government through the Norwegian MoD and articulated through long-term plans. When the policy was first put into effect in the 2000s, the MoD was also the commissioning body for acquisition projects. This meant they were responsible for initiating acquisition projects and providing frameworks and guidelines for implementation, based on overarching policies and strategies. The Armed Forces had the role of planner and executor of acquisitions, within MoD's given frameworks and guidelines. At first glance, this seems to be a straightforward process with a limited number of actors. In reality, various parts of the MoD were responsible for issuing the new policy and for launching projects. Multiple actors within the Armed Forces were also involved, ranging from end-users and engineering divisions to contract management. Although the overarching policy stated that off-the-shelf should be the preferred acquisition strategy, there were several decision points where the policy could be overridden by other, more or less specific, goals and guidelines in the projects, such as industrial policy requirements or operational requirements specifications (Presterud et al., 2016). Can a simplification of the process and governance structure in recent years be aligned with the observed development in the implementation of the policy?

As in previous years, in 2023 the MoD continued to both articulate the policy through the long-term plans and act as the commissioning authority for acquisition projects. However, by this time, their role as commissioning authority was limited to only the largest projects. Smaller projects were initiated by the Armed Forces through the Defense Staff. All projects were now required to undergo a conceptual study, a requirement that had previously applied only to the largest projects. Responsibility for the conceptual study rested with the Armed Forces, while further analysis and development of the decision basis, as well as the planning and implementation of the projects, were now handled by the Norwegian Defense Materiel Agency<sup>3</sup>. As before, several different actors remained involved within each of the different organizations.

Although we are unable to specify exactly how many actors were involved in decisions related to using off-the-shelf solutions as an acquisition strategy in either 2014 or 2023, there is little evidence to suggest that the process and decision-making hierarchy have become less complex over time. We can largely conclude that simplifications in processes and governance structures do not align with the observed strengthening of the off-the-shelf policy's implementation in the years following 2020.

Finally, within the category of structural conditions, we examine the incentives for adopting the policy. The literature asserts that policies which entail a cost or burden for those tasked with implementation must be accompanied by incentives to encourage actors to comply.

---

<sup>3</sup> Established in 2016 by separating the part of the Norwegian Armed Forces' Logistics Organization that was responsible for acquisition.

The off-the-shelf policy cannot be characterized as such a “costly” policy. The policy essentially requires actors to choose one acquisition strategy over another. In itself, this is not costly, in fact, conducting off-the-shelf procurements is less resource-intensive than engaging in demanding development processes (Presterud et al., 2016). For individuals, however, it can still be experienced as a burden to change behavior, that is, to do things differently. It may also be perceived as a sacrifice or cost to move away from exciting development projects in favor of simpler off-the-shelf procurements. A 2015 study of incentives in the Norwegian defense sector’s procurement process found that actors used their expert authority to steer projects toward technologically sophisticated, development-oriented solutions, partly to serve their own interests (Presterud & Øhrn, 2015). In the 2023 survey, 27% responded that they found projects involving development most inspiring, while only 6% preferred projects without development. It is therefore conceivable that some formal incentives may be necessary to encourage actors to opt for more off-the-shelf solutions.

In the Norwegian public sector, there is little tradition for individual incentive schemes, either in the form of rewards or sanctions. Instead, government agencies are primarily governed through reporting and management objectives, which create incentives at the organizational or group level. When the off-the-shelf policy was first introduced, however, it was not accompanied by such objectives or reporting requirements. A performance indicator related to the share of off-the-shelf procurements was not introduced until 2016 but this indicator was discontinued in 2020 (Forsvarsdepartementet, 2016b).

The literature suggests that the observed increase in off-the-shelf procurements in 2023 may be due to changes in incentive schemes for the implementation of the policy. However, this development does not correspond with the observed changes in incentive schemes over time. Thus, it is unlikely that the development in the implementation of the off-the-shelf policy in acquisition projects is due to incentive arrangements.

In this section, we have examined the structural conditions related to the off-the-shelf policy, to assess whether changes in these have influenced its implementation. Our analysis is based on publicly available documents and may therefore not capture all relevant communication or informal guidance and practices. Within these limitations, our review of communication, process, management structure, and incentive schemes indicates limited alignment between strengthened structural conditions and the observed development in implementation of the policy.

### **Non-structural Conditions**

Defense acquisition is not a mechanical process, but a process carried out by people who make their assessments and decisions in various ways based on their inherent knowledge and perspectives. How the implementing actors respond to policy change is therefore crucial for its outcome. In this section, we examine whether changes among the implementing actors align with the observed development in implementation of the off-the-shelf policy in acquisition projects.

We have limited historical data on actors’ attitudes toward the off-the-shelf policy. Still, we know a few things: In 2013, the then Chief of Defense in Norway, Harald Sunde, commented on off-the-shelf acquisitions (Bentzrød, 2013). In his statement, he said that Norway should choose development projects and not buy off-the-shelf materiel when “expensive and heavy weapon systems” are to be acquired. In our 2014 survey, we asked actors in the acquisition workforce if they agreed with the Chief of Defense on this. Around 50% disagreed with him, while about 25% agreed. The remaining 25% were neither in agreement nor disagreement.

Half of those asked disagreed that off-the-shelf is the wrong solution for “large” acquisitions, which can be interpreted as them having a positive attitude toward the possibilities



of using off-the-shelf as an acquisition strategy. But there were also half who either had no opinion or who agreed with the Chief of Defense. In a study from 2015 we found that actors perceived off-the-shelf solutions as something that can only work as an acquisition strategy in special cases (Presterud et al., 2016). Here, the actors mainly pointed to *one* argument for choosing off-the-shelf solutions: it can be procured quickly. At the same time, we found that acquisition speed was not a priority (Presterud & Øhrn, 2015). Given these findings, it is perhaps not surprising that the policy was relatively rarely exercised in 2015.

In 2023, our second survey was conducted. This shows a clear development in the actors' attitudes towards off-the-shelf solutions. In 2023, the proportion of respondents who disagreed with the former Chief of Defense had increased to 73%, while only 7% now agreed with him. Although the proportion who agreed with the then Chief of Defense in 2014 was already quite low, this is a significant shift. A full 71% of the actors state in the same survey that the focus on procuring and using off-the-shelf solutions in the Armed Forces was increasing, and only 16% disagreed. This was despite the absence of any changes in the structural conditions surrounding the policy. Only 6% of respondents stated that relevant off-the-shelf solutions are rarely available when the defense sector procures materiel. In 2023, implementing actors therefore appear to be both well acquainted with the policy and largely confident in its practical feasibility.

It appears that a change has occurred among the actors over the years, even though the way decision-makers govern the policy has not changed significantly. This suggests that non-structural factors provide a more plausible interpretation of the observed change than structural measures introduced by decision-makers. This aligns with other findings from the 2023 survey, where only 9% state that pressure from the MoD has led them to make different choices than they otherwise would have when it comes to implementing the policy.

We find indications of a change in attitudes among the actors, and these coincide with the observed changes in implementation of the off-the-shelf policy in acquisition projects. However, the change in attitudes does not coincide with corresponding changes in how the policy has been governed through structural conditions. This raises the question of what has changed among the implementing actors. Because of data limitations and privacy considerations, we can only examine a limited set of characteristics. One relevant factor highlighted in the literature is experience.

We find that more experienced individuals are both more likely to disagree with the former Chief of Defense and more likely to believe that off-the-shelf solutions are available.<sup>4</sup> This suggests that experience is associated with attitudes toward the policy. However, we cannot assess whether experience has driven implementation in practice, because we cannot link individual respondents to decisions made in specific acquisition projects.

Is it likely that an overall increase in experience among the acquisition workforce can help explain the positive development in implementation of the policy? Survey respondents are, on average, more experienced in 2024 than in 2014, but we cannot assess whether the workforce experience level has increased across the acquisition system as a whole. Available key figures from the Norwegian Defence Materiel Agency do not indicate a major increase, and we lack comparable data for the other agencies involved (Forsvarsmateriell, 2017, 2024; Statistisk Sentralbyrå, 2025a, 2025c, 2025b).

So, even though experience appears to influence the actors' attitudes toward the policy, there is little to suggest that the increase in the use of off-the-shelf as an acquisition strategy is

---

<sup>4</sup> We define experience as how many years the actors have worked with procurement in the defense sector, divided into three categories (1–3 years, 4–6 years, and over 6 years)

due to major changes in the experience level of the actors. In fact, the survey shows that all subgroups have developed more positive attitudes toward the policy over time. The inexperienced group has actually changed the most, even though they are still more skeptical than the experienced group. As far as we know, there have not been any other significant changes in the type of people working in the acquisition workforce during this period that could help explain the shift.

Another possible explanatory factor is the actors' sense of meaningfulness. The implementing actors' understanding of how a policy contributes to important societal goals influences their willingness to implement it (Tummers et al., 2012). When guidelines for using off-the-shelf as an acquisition strategy were originally formulated, the need for it was primarily justified by the necessity to save money. We do not have direct measures of how meaningful implementing actors perceived this goal to be, but we do know that the Armed Forces viewed the general cost-efficiency requirements imposed on them during this period as very demanding (Forsvarssjefen, 2015).

In the years that followed, there was a significant shift in the geopolitical situation and in Norway's security policy focus. Flat or reduced defense budget trajectories were replaced by renewed investments in the Armed Forces and increased budgets (Forsvarsdepartementet, 2020, 2023b). Along with this came a shift in the argumentation for the off-the-shelf policy. Even though the argument about cost efficiency was still mentioned in the governing documents, there was now clearly increased focus on off-the-shelf as a strategy intended to speed up acquisitions.

Here, we observe a policy that has evolved over time. From being promoted primarily as a cost-efficiency measure, it is now promoted as a policy whose main goal is to build up the Armed Forces as quickly as possible. While we cannot empirically test whether this shift in justification has changed actors' sense of meaningfulness, the mechanism is consistent with earlier studies showing that implementing actors viewed off-the-shelf solutions as particularly suitable when speed is prioritized (Presterud et al., 2016). Given the increased emphasis on rapid strengthening of the Armed Forces, it is therefore plausible that the policy is perceived as more meaningful when framed as a means to accelerate capability build-up than when framed primarily as a cost-efficiency measure.

Overall, our analysis indicates that changes among implementing actors, particularly their attitudes and perceptions of the off-the-shelf policy, coincide with the increased share of off-the-shelf acquisition projects observed after 2019. While the available data do not support strict causal claims, alternative explanations based on tractability and formal structural conditions fit the empirical patterns less well. We cannot determine precisely why implementers' attitudes and perceptions shifted, but we outline one plausible mechanism: as the policy justification has shifted from primarily cost efficiency toward speed and the ambition of more rapid operational effect in a changed security environment, the policy may have become more salient and meaningful to implementers, increasing their willingness to select off-the-shelf solutions at key decision points.

## **Conclusion**

Over the past ten years, the security policy context has changed dramatically. Like many other countries, Norway has recognized the need for a rapid buildup of national defense capabilities. To achieve this, the acquisition sector must be further developed. However, acquisition reform has proven to be difficult and time-consuming to implement.

The purpose of this study has been to examine which conditions have shaped the implementation of the policy promoting increased use of off-the-shelf solutions within Norway's defense acquisition system. The policy was promoted as early as in the 2000s, but for a long



time seemed to have negligible effect. It is only in the years after 2020 that we observe a strengthening in the implementation of the policy. Based on policy implementation theory, we propose three hypotheses about overarching conditions that may have shaped this development. We explore the hypotheses using empirical data from document analysis, previous research, and data from surveys. The three hypotheses serve different analytical purposes. The tractability hypothesis primarily functions as a contextual plausibility check, while the latter two hypotheses capture factors more directly related to implementation dynamics.

The data do not allow for strict causal identification, but our analysis provides a theory-informed interpretation of the observed patterns of implementation over time. Table 2 summarizes the hypotheses derived from the literature and the degree of empirical alignment identified in the analysis.

**Table 6. Hypotheses and Findings**

| Factors                   | Hypotheses                                                                                                                                                           | Findings           |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| Tractability              | The observed development in implementation of the off-the-shelf policy in acquisition projects may be related to the problem having become more tractable over time. | Weak alignment     |
| Structural conditions     | The observed development in implementation of the off-the-shelf policy in acquisition projects may be attributable to strengthened structural conditions.            | Weak alignment     |
| Non-structural conditions | The observed development in implementation of the off-the-shelf policy in acquisition projects may be attributable to changes among the implementing actors.         | Stronger alignment |

The three hypotheses receive different levels of support. The analysis points to limited alignment between the implementation shift, measured as the increased share of off-the-shelf acquisition projects, and changes in external conditions (tractability). While the defense market has evolved in recent years, the available information does not allow us to determine whether procurement of off-the-shelf solutions has become systematically easier or harder. Moreover, the Norwegian acquisition portfolio does not appear to have become more “off-the-shelf-friendly” over the period. We also find little support for the idea that the observed development primarily reflects the passage of time.

The available evidence also aligns weakly with explanations emphasizing changes in formal structural conditions. There have been some changes in communication, processes, governance structures, and incentives during the period, but these changes do not align with the development we observe in off-the-shelf acquisitions during the same time.

By contrast, explanations focusing on non-structural conditions, implementing actors’ attitudes and perceptions, provide the most consistent interpretation of the observed development. What has driven these changes remains uncertain, as neither the governance structure nor the composition of the group of implementing actors appears to have changed substantially. One plausible explanation is that the policy is now perceived as more meaningful because it is increasingly justified by the need for rapid build-up of the Armed Forces rather than primarily by cost efficiency.

Overall, the study’s findings indicate that factors outside direct control of decision-makers can impact whether a policy actually succeeds. For decision-makers, this may seem unsettling. However, this can also be seen as something positive because it means that real change does not necessarily require major structural adjustments. Policies linked to meaningful





goals that actors perceive as significant, in this case the need for rapid build-up of defense capability, can mobilize implementation willingness among the implementing actors and support the desired development.

Furthermore, the study illustrates that implementation can take time. This implies that repeated reforms and frequent adjustments in governance signals can weaken the possibility of both achieving effects and learning. For a sector facing historic demands to deliver faster, paradoxically, many structural reforms can contribute to delaying the necessary development. A more knowledge-based and evaluation-driven approach may be necessary to ensure that reform measures actually have the time and opportunity to take effect.

Overall, the study suggests that further reform efforts should combine clear political priorities with long-term stability, while also giving greater emphasis to the motivation, knowledge, and sense of purpose of the implementing actors. Effective acquisition reforms and policies are not just about what is decided, but equally about how those decisions are perceived, owned, and put into practice by the people who actually carry them out.

## Sources

- Bentzrød, S. B. (2013, November 2). *Norge må kjøpe uferdige milliardvåpen* [Nettavis]. Aftenposten.no. <https://www.aftenposten.no/norge/i/m6dng/norge-maa-kjoepe-uferdige-milliardvaapen>
- Berg, H., & Waage, K. (2021). *Effektive materiellanskaffelser i Forsvaret—Øker andelen hyllevarekjøp?* (FFI-rapport No. 20/03147).
- Forsvarsdepartementet. (2004a). *Den videre moderniseringen av Forsvaret i perioden 2005–2008—IVERKSETTINGSBREV FOR FORSVARSSEKTOREN*.
- Forsvarsdepartementet. (2004b). *St.prp. Nr. 42 (2003–2004)—Den videre moderniseringen av Forsvaret i perioden 2005–2008*.
- Forsvarsdepartementet. (2008a). *Et forsvar til vern om Norges sikkerhet, interesser og verdier: Iverksettingsbrev til forsvarssektoren (2009–2012)*.
- Forsvarsdepartementet. (2008b). *Prop. 48 (2007–2008) Et forsvar til vern om Norges sikkerhet, interesser og verdier*.
- Forsvarsdepartementet. (2012a). *Prop. 73 S (2011–2012) Et forsvar for vår tid*.
- Forsvarsdepartementet. (2012b, June 28). *Et forsvar for vår tid—Iverksettingsbrev til forsvarssektoren for langtidsplanperioden 2013–2016*.
- Forsvarsdepartementet. (2014). *Framtidige anskaffelser til Forsvaret (FAF) 2014–2022*.
- Forsvarsdepartementet. (2015). *Rapport om anbefalinger knyttet til tiltak utarbeidet av McKinsey & Company Inc.*
- Forsvarsdepartementet. (2016a). *Prop. 151 S (2015–2016) Kampkraft og bærekraft*.
- Forsvarsdepartementet. (2016b, December 20). *IVB LTP 2017–2020 VEDLEGG A2 – MÅL-, RESULTAT- OG RISIKOSTYRING FORSVARSMATERIELL*.
- Forsvarsdepartementet. (2020). *Prop. 62 S (2019–2020)—Vilje til beredskap – evne til forsvar Langtidsplan for forsvarssektoren*.
- Forsvarsdepartementet. (2023a). *Framtidige anskaffelser til forsvarssektoren (FAF) 2023–2030*.



- Forsvarsdepartementet. (2023b, October 6). *Statsbudsjettet 2024: Et styrket forsvarsbudsjett i en krevende tid* [Nyhet]. regjeringen.no. <https://www.regjeringen.no/no/aktuelt/forsvarsbudsjettet-okker-med-20-prosent/id2997543/>
- GAO. (2024). *Weapon systems annual assessment: DoD is not yet well-positioned to field systems with speed* (Report to Congressional Committees GAO-24-106831).
- GAO. (2025). *Defense acquisition reform: Persistent challenges require new iterative approaches to delivering capability with speed* (GAO-25-108528).
- Giacchino, S., & Kakabadse, A. (2003). Successful policy implementation: The route to building self-confident government. *International Review of Administrative Sciences*, 69(2), 139–160. <https://doi.org/10.1177/0020852303069002002>
- Hill, M., & Hupe, P. (2003). The multi-layer problem in implementation research. *Public Management Review*, 5(4), 471–490. <https://doi.org/10.1080/1471903032000178545>
- House of Commons Defence Committee. (2023). *It is broke—And it's time to fix it. The UK's defence procurement system* (Ninth Report of Session 2022–23).
- Hupe, P. L. (2011). The thesis of incongruent implementation: Revisiting Pressman and Wildavsky. *Public Policy and Administration*, 26(1), 63–80. <https://doi.org/10.1177/0952076710367717>
- Liang, X., Tian, N., Lopes Da Silva, D., Scarazzato, L., Karim, Z. A., & Guiberteau Ricard, J. (2025). *Trends in world military expenditure, 2024*. Stockholm International Peace Research Institute. <https://doi.org/10.55163/AVEC8366>
- Matland, R. E. (1995). Synthesizing the implementation literature: The ambiguity-conflict model of policy implementation. *Journal of Public Administration Research and Theory: J-PART*, 5(2), 145–174.
- McLaughlin, M. W. (1987). Learning from experience: Lessons from policy implementation. *Educational Evaluation and Policy Analysis*, 9(2), 171–178. <https://doi.org/10.2307/1163728>
- National Audit Office. (2021). *Improving the performance of major equipment contracts* (SESSION 2021-22 24 JUNE 2021 HC 298).
- O'Toole, L. J. (2000). Research on policy implementation: Assessment and prospects. *Journal of Public Administration Research and Theory*, 10(2), 263–288. <https://doi.org/10.1093/oxfordjournals.jpart.a024270>
- Prebensen, F., & Voldhaug, J. E. (2026). *Materiellinvesteringer i forsvarssektoren – årsaker til forsinkelser* (FFI-Rapport No. 26/004).
- Pressman, J. L., & Wildavsky, A. B. (1973). *Implementation: How great expectations in Washington are dashed in Oakland*. University of California Press.
- Presterud, A. O., Øhrn, M., & Berge, H. (2016). *Effektive materiellanskaffelser i Forsvaret—Økonomiske gevinster ved økte hyllewareanskaffelser* (FFI-rapport No. 2015/02332). <https://www.ffi.no/publikasjoner/arkiv/effektive-materiellanskaffelser-i-forsvaret-okonomiske-gevinster-ved-okte-hyllewareanskaffelser>
- Presterud, A., & Øhrn, M. (2015). *Effektive materiellanskaffelser i Forsvaret- en studie av insentiver i investeringsprosessen* (FFI-rapport No. 2015/00555).



- Riksrevisjonen. (2005). *Riksrevisjonens undersøkelse av materiellinvesteringsprosjekter i Forsvaret* (Dokument nr. 3:16 (2004–2005)). Riksrevisjonen.
- Riksrevisjonen. (2024). *Riksrevisjonens undersøkelse av gjennomføring av materiellinvesteringer i forsvarssektoren* (Dokument 3:16 (2023–2024)). Riksrevisjonen.
- Sabatier, P., & Mazmanian, D. (1979). The conditions of effective implementation: A guide to accomplishing policy objectives. *Policy Analysis*, 5(4), 481–504.
- Sabatier, P., & Mazmanian, D. (1980). The implementation of public policy: A framework of analysis. *Policy Studies Journal*, 8(4), 538–560. <https://doi.org/10.1111/j.1541-0072.1980.tb01266.x>
- Sætren, H. (2024). *Chapter 1: An intellectual history and state-of-the-art assessment of policy implementation research and theory*. <https://www.elgaronline.com/edcollchap/book/9781800885905/book-part-9781800885905-8.xml>
- Sager, F., Mavrot, C., & Keiser, L. R. (2024). Introduction: The notion of policy implementation and why it is important. In F. Sager, C. Mavrot, & L. R. Keiser (Eds.), *Handbook of public policy implementation* (pp. 1–10). Edward Elgar Publishing. <https://doi.org/10.4337/9781800885905.00006>
- Scarazzato, L., Tian, N., Lopes Da Silva, D., Liang, X., Karim, Z. A., & Guiberteau Ricard, J. (2025). *The SIPRI top 100 arms-producing and military services companies, 2024*. Stockholm International Peace Research Institute. <https://doi.org/10.55163/XPGD6816>
- Tummers, L., Steijn, B., & Bekkers, V. (2012). Explaining the willingness of public professionals to implement public policies. *Public Administration*, 90, 716–736. <https://doi.org/10.1111/j.1467-9299.2011.02016.x>
- Van Meter, D. S., & Van Horn, C. E. (1975). The policy implementation process: A conceptual framework. *Administration & Society*, 6(4), 445–488. <https://doi.org/10.1177/009539977500600404>
- Voldhaug, J. E., Prebensen, F., & Presterud, A. O. (2024). *Materiellinvesteringer i forsvarssektoren – når vi målene?* (FFI-rapport No. 24/00554).



# Going on the Offensive: Acquisition for Competing with China

**Emily de La Bruyere**—is a co-founder of Horizon Advisory and senior fellow at the Foundation for Defense of Democracies. She holds a BA summa cum laude from Princeton University and an MA summa cum laude from Sciences Po, Paris.

**Nathan Picarsic**—is a co-founder of Horizon Advisory and senior fellow at the Foundation for Defense of Democracies. He holds a BA from Harvard College and has completed executive education programs at Harvard Business School and Defense Acquisition University.

## Abstract

This paper argues that the great power competition with China should drive defense acquisition processes, and that those processes should be used proactively to compete with Beijing, not simply as reactive levers. This applies to the directions in which the United States invests: The United States should prioritize capabilities not only that defend against PRC capabilities but also that have the potential to generate desired effects on China. This proactive logic also applies to the way in which the United States times and signals its resource allocations. The US defense acquisition system should invest and message in a way that accounts for and targets the PRC's system – to ensure that US defense acquisition processes not only support American positioning and perceptions of it, but also appropriately shape the adversary. Across the board, an adversary-informed, proactive acquisition system should begin from an understanding of China's central planning system, and use that system both to identify PRC priorities, capabilities, and investments early and to time US responses to maximize strategic effect.

## Introduction

China poses a near-peer – and, increasingly, peer – threat to the United States, across economic, science and technology, and military vectors. This makes China the pacing threat for the United States; the dominant force driving adversary-informed US positioning. The challenge posed by the People's Republic of China (PRC) shapes US concerns about and investments in everything from supply chain security to scientific and technological innovation. In particular, China's challenge shapes the US defense apparatus, including capability development, capability demonstration, and the defense acquisition system more broadly.

But thus far, US efforts to counter China's threat – militarily, socially, economically, and industrially – have tended to be defensive and reactive. The US has sought to develop alternative sources of supply where China controls value chains, rather than to cultivate positions of leverage vis-à-vis Beijing. In the defense sector, the US has focused defensively on bolstering capabilities threatened by China and reactively on keeping pace with Chinese investments. The reaction to Chinese restrictions on critical minerals, like gallium and germanium, and the subsequent rush to secure rare earth element and REE permanent magnet supply lines proves the point acutely and recently.

This paper argues that the long-term, peacetime competition with China should drive defense acquisition processes. Moreover, those processes should be used proactively to take the offensive, and to compete with Beijing, not simply as reactive levers. This applies to the directions in which the United States invests: The United States should prioritize capabilities not only that defend against PRC capabilities but also that have the potential to generate desired effects on China. This proactive logic also applies to the way in which the United States times and signals its resource allocations. The US defense acquisition system should invest and message in a way that accounts for and targets the PRC's system – to ensure that US defense acquisition processes not only support American positioning and perceptions of it, but also



appropriately shape the adversary.

To do this – to adopt a proactive, adversary-informed defense acquisition process – requires understanding and monitoring China’s present capabilities and how it intends for those to evolve, where it is investing, what its priorities in those fields are, and how they are paying off. The United States must also understand and monitor the underlying PRC systems: the systems according to which China’s priorities are formed, investments are determined, and capabilities are evaluated – and also and especially how those systems respond to external stimuli (e.g., the US acquisition processes). Then, the United States can direct its acquisition process accordingly; the United States can incorporate threat intelligence further to the left in the acquisition process and on a recurring basis across research and development (R&D), procurement, and operations and maintenance investments.

An enhanced understanding of Chinese systems can begin with the country’s high-level planning apparatus: the marquee National Five-Year Plans for Economic and Social Development (hereafter National Five-Year Plans), the subordinate plans that they inform, and the resources that those in turn drive; the cycles according to which this process operates; and the degrees to and ways in which all respond to external stimuli. China’s Party and state leaders oversee a “unified planning system” (Dong, 2025). That system starts, at the highest level, with a national architecture that develops National Five-Year Plans. Those National Five-Year Plans in turn inform an ecosystem of subordinate plans, including National-Level Special Plans; industry-, provincial-, and locality-specific plans; and project- and task-specific implementation guides.

A distinct, enduring, and structured feature of Chinese governance, this system has been executed since the country’s founding – and is increasingly institutionalized. The development of the central Five-Year Plan follows a prescribed sequence of steps, is undertaken by a prescribed cast of stakeholders, and generates a prescribed set of outputs. That central Five-Year Plan in turn catalyzes development of a prescribed system of subordinate plans (e.g., National-Level Special Plans) and determines set categories of resource allocations. This makes the planning process both a ready analytical target for monitoring the PRC and a ready context for efforts to compete with China.

The PRC’s hypersonic program offers an example both of the role of the planning system in shaping PRC capability development and the analytical utility of tracking that system. The 2006 11th National Five-Year Plan established – and subordinate plans elaborated on and operationalized – a set of “Major Science and Technology Projects,” slated to last through 2020. One of those focused on hypersonic capabilities (Guancha, 2022). The project supported development of, among other things, China’s DF-ZF hypersonic glide missile and Dongfeng-17 medium-range ballistic missile. The PRC first tested the DF-ZF in 2014 and the Dongfeng-17 in 2017 (Science and Technology Daily, 2018).

Those tests, and the continued capability development and demonstration that have followed, appear to have taken the United States by surprise. As the Financial Times reported of a 2021 PRC hypersonic test, “China tested a nuclear-capable hypersonic missile in August, showing a capability that caught U.S. intelligence by surprise” (Reuters, 2023). And in 2022, speaking about the Dongfeng-17, Dr. Mark Lewis, executive director of the National Defense Industrial Association’s Emerging Technologies Institute, and former director of defense research and engineering for modernization at the Pentagon, stated that, “the investments that we’ve seen the Chinese make in hypersonics are frankly startling” (Harper, 2022). Moreover, taken by surprise, the United States has responded defensively and reactively to the PRC’s hypersonics program, focusing on mirroring and defending against Chinese capabilities.





But close monitoring of the PRC's planning apparatus could have provided early indications of the PRC's ambitions, corresponding investments, and operationalizing entities. In addition, China's planning apparatus responds to external stimuli: The United States could proactively leverage the US defense acquisition process to shape PRC ambitions, investments, and systems – rather than being shaped by the PRC – if better able to incorporate sensing of Chinese priorities and capability developments in the earliest stages of US acquisition processes.

This paper begins with a survey of the PRC's high-level planning system. The next section uses the 2006 to 2020 National Science and Technology Major Projects to detail ways in which the Chinese planning system shapes resource allocations and responds to external stimuli. The third section details the insight into PRC priorities, capabilities, and entities that can be gleaned from understanding and monitoring the planning system. And the fourth section discusses ways in which the United States might leverage understanding of the Chinese governance system to develop not only an adversary-informed but a proactive, adversary-informed defense acquisition process, able to shape the PRC in a period of long-term, peacetime competition.

## **Understanding the Planning System: The Five-Year Plan Cycle**

China's Five-Year Plans constitute the most important blueprints for the country's strategic objectives. These plans detail national priorities and goals and outline measures through which to pursue them. Five-Year Plans exist at the strategic level. They establish roadmaps for the PRC over a relatively long time period. They do not detail specific policies or funding streams.

The strategic blueprints outlined in Five-Year Plans are translated into more concrete lines of effort – as well as allocations of responsibility and resources – through a system of subordinate plans. These include National-Level Special Plans, of which there are dozens per Five-Year Plan cycle, as well as regional and institutional Five-Year Plans. National-Level Special Plans are particularly important within this ecosystem of subordinate plans. They operationalize the major initiatives detailed in National Five-Year Plans, including by detailing their more specific goals, requirements, and indicators; assigning them to bureaucratic stakeholders for oversight; and, therefore, connecting them to resourcing (Yao, 2012).

All of these plans are developed through a regimented, multi-step process. The process for formulating the PRC's high-level, central Five-Year Plan begins halfway through the previous Five-Year Plan period. During the third year of that period, the National Development and Reform Commission (NDRC) organizes a mid-term evaluation of it. The evaluation generates a report that is approved by the State Council and submitted to the National People's Congress. The evaluation also informs research for the next Five-Year Plan. That research takes the form of reports, commissioned by the NDRC. In the third step of the planning process, drafting, the NDRC distills the results of the research phase into a draft plan, completed by the end of the fourth year of the previous Five-Year Plan period (China Daily, 2025).

The Central Committee of the Chinese Communist Party (CCP) revises the NDRC's draft in the final year of the previous Five-Year Plan – with participation from representatives from the National People's Congress, Chinese People's Political Consultative Conference, State Council, and local governments and in consultation with external stakeholders. The Central Committee's suggestions are then passed, after which point the NDRC reconciles them with its framework for the new Five-Year Plan, and then solicits opinions from government and non-government entities. Around the start of the first quarter of the new Five-Year Plan period, the draft plan is approved and released.



Subordinate plans are developed according to similar, if less laborious, processes. For example, PRC sources summarize the steps for formulating National-Level Special Plans as "1. democratic participation, 2. planning coordination, 3. expert consultation, 4. approval and publication, and 5. evaluation and revision" (Yang, 2011). Generally, the sequence is initiated with a proposal for the plan submitted to the State Council's Central Committee. The State Council then deliberates over submitted proposals, after which point Planning Coordination Groups are created for approved plans. The groups begin by developing a work plan and methodology, before soliciting opinions, conducting research, and drafting a plan. Next, they revise and improve the initial draft with input from relevant State Council Departments and prominent external entities (e.g., universities). After a revised draft is socialized through "public consultation," final revisions are made and the plan is reviewed at an executive meeting of the State Council (Yang, 2011).

## **Understanding the Planning System: The National Science and Technology Major Projects**

The PRC's 11th Five-Year Plan system established 16 National Science and Technology Major Projects (国家科技重大专项) in 2006, slating those to last through 2020. These projects offer an example both of the role of the planning system in shaping PRC capability development, including military capability, and the analytical utility of tracking that system. The PRC published the 11th National Five-Year Plan in 2006, outlining the country's strategic direction for the 2006 to 2010 period. That plan announced that China would "launch a number of Major Science and Technology Projects."

The 11th Five-Year Plan provided little additional detail. But the responsive, Medium- and Long-term Development Plan for Science and Technology Development (国家中长期科学和技术发展规划纲要) (2006-2020) (MLDP) did. The MLDP, a National-Level Special Plan, constituted a 15-year, strategic roadmap for China's scientific and technological ambitions, programs, and priorities. And the policy documented elaborated on the Major Projects at which the National Five-Year Plan had hinted, explaining that these would be national research, development, and industrialization initiatives. "Starting from pressing real-world needs, focusing on breakthroughs in major key and common technologies," the projects would "achieve leapfrog development and fill gaps." The MLDP stated that there would be 16 Major Projects. It listed 13, ranging from "Ultra-Large-Scale Integrated Circuit Manufacturing Technology and Complete Processes" to "High-Resolution Earth Observation Systems." The remaining three were classified and remained undisclosed in PRC formal media, policies and plans, and resourcing disclosures.

After and in response to the MLDP's publication, subsequent layers of the PRC planning apparatus continued to operationalize and provide detail on the Major Projects. The MLDP described the ambitions guiding the establishment of the Major Projects. But it did not elaborate on any of them beyond their high-level names. Nor did it indicate related tasks, designate lead units, or identify performance metrics. Those details were addressed in another, paired National-Level Special Plan, the 11th Five-Year Plan for Science and Technology Development ("十一五"科学技术发展规划) – as well as in the more detailed planning and bureaucratic work that flowed from it.

Where the MLDP introduced the Major Projects, the 11th Five-Year Plan for Science and Technology Development initiated their implementation. Published in August 2006, this plan identified, as its first "key task," "implementation" of the MLDP under the strategic framework of the 11th Five-Year Plan for Economic and Social Development. Accordingly, the 11th Five-Year



Plan for Science and Technology Development detailed both the “development strategy” behind the Major Projects and, for the 13 disclosed ones, goals for the 2006 to 2010 period (i.e., their objectives for the first Five-Year Plan cycle of their period of performance). The plan also clarified that the National Science and Technology Education Leading Group was the authority charged with “implementation leadership” for the Major Projects.

Bureaucratically, the National Science and Technology Education Leading Group was positioned in the General Office of the State Council.<sup>1</sup> Its members included the Director of the National Development and Reform Commission, the Minister of Education, the Minister of Science and Technology, the Minister of Industry and Information Technology, the Minister of Finance, the Minister of Agriculture, the President of the Chinese Academy of Sciences, the President of the Chinese Academy of Engineering, the Deputy Secretary-General of the State Council, the Director of the National Natural Science Foundation of China, and the Executive Vice President of the China Association for Science and Technology (State Council, 2008).

Based on the authority granted it by the PRC planning process, the National Science and Technology Education Leading Group assigned the Ministry of Science and Technology (MOST) the role of operationalizing the Major Projects.<sup>2</sup> MOST proceeded to formulate management regulations for the initiatives and to assign them to “leading organizations.”

“Leading organizations” refer to the entity or entities responsible for specific organization and implementation of any given Major Project. Most of the 16 MLDP-launched Major Projects were assigned to government departments or ministries (e.g., the Ministry of Industry and Information Technology). But local governments, State-owned enterprises, and academies were also eligible. The leading organizations established management offices to undertake day to day work on their Major Projects; assembled overall expert groups, composed of individual experts, to consult on technical trajectories; and recruited participants, including academic, industry, and government research entities, often working together, to carry out project tasks. The leading organizations also formulated implementation plans for their Major Projects. Projects were formally launched only after their implementation plans were approved by the State Council.

This process can present as self-contained and insulated. But it featured engagement with external stakeholders and responded to external stimulus. On the first front, implementation plans were developed through consultation with researchers, management experts, financial experts, enterprises, and local government stakeholders (Unitalen, 2026). On the second front, some Major Projects were selected in response to external pressures and others were accelerated because of external pressures. One of the MLDP-era Major Projects, for instance, focused on “large aircraft” development; PRC sources attribute that project to Israel’s 2000 decision to terminate cooperation with China on early warning aircraft (Shanghai Commission, 2018). Another Major Project – the Core Electronic Devices, High-End General-Purpose Chips, and Basic Software Products Major Project – sought to localize chips, software, and electronic devices.<sup>3</sup> PRC sources suggest that the Major Project’s 2008 launch was accelerated by that year’s “Microsoft black screen incident,” Microsoft’s forced black screen warnings to users of

---

<sup>1</sup> This State Council affiliated bureaucratic organ would see multiple evolutions in the time since its founding in 2003: It was reorganized as the National Leading Group for Science and Technology in 2018; in 2023, its primary functions, including “studying and reviewing national science and technology development strategies and plans, major science and technology tasks, and coordinating cross-departmental science and technology matters,” were transferred to the Central Science and Technology Commission.

<sup>2</sup> Or, at least, the civilian Major Projects. The seven military Major Projects appear to have been assigned to the General Armaments Department of the People’s Liberation Army.

<sup>3</sup> The Major Project is often referred to as “Special Project 01,” because it was first on the MLDP’s list of Major Projects.



pirated Windows and Office (Science Times, 2008). The PRC viewed this as an information security threat – and, through the Major Project, prioritized development of an independent Linux operating system and basic office software.

Importantly, the Major Projects – with them the PRC’s national planning system as mapped in this example – drove military as well as civilian capability development. One of the three classified Major Projects, for example, focused on hypersonic vehicle technology. The project incubated PRC hypersonic capabilities – including the DF-ZF hypersonic glide missile and Dongfeng-17 medium-range ballistic missile – that, when demonstrated, took the United States by surprise (Science and Technology Daily, 2018).

Finally, while the Major Projects outlined in the MLDP were slated to operate only through 2020, the Major Project program remains active. In 2016, at the start of the 13th Five-Year Plan period, the PRC introduced a new generation of Major Projects, the Innovation 2030 Major Projects. Some of those are obvious successors to the MLDP era set. For instance, where the C919 and Y-20 developed under the Large Aircraft Major Project relied on foreign engines, the Innovation 2030 Major Projects include one dedicated to localizing engine technology, the Aeroengine and Gas Turbine Major Project. Others are new and respond to emergent technological trends. The New Generation Artificial Intelligence Major Project constitutes an obvious example. Many of the Innovation 2030 Major Projects have military relevance. All continue to reflect some of the PRC’s highest-level, most capital-intensive research, development, and industrialization efforts.

## **Monitoring the Planning System to Monitor the PRC**

The PRC’s planning system is responsible for, at the highest level, defining the country’s strategic vision, and, at a more granular level, obligating state funds to pursue that vision. Monitoring the planning system therefore enables monitoring, contextualization, and anticipation of China’s strategic priorities, operational activities, and capability development. And doing so is critical for an adversary-informed defense acquisition process.

In other words, early identification of the priorities established in the planning system provides early warning of the capabilities Beijing is positioned to develop. Monitoring of the process through which those priorities are operationalized fuels contextualized, adversary-informed insight into the PRC’s progress in developing those capabilities. And mapping of the entities involved (e.g., assigned to lead specific initiatives) permits targeting of linchpin adversarial actors. Such early warning, adversary-informed insight, and targeting can all support adversary-informed defense acquisition that accounts for PRC goals and trajectories before they are proven – as well as strengths and weaknesses that might not otherwise be known.

Again, the Major Projects offer a prime case. The PRC stated its highest-level, most forward-looking scientific and technological goals in 2006, when it established the 16 National Science and Technology Major Projects. Those projects made clear that the PRC would, between 2006 and 2020, allocate significant government resources toward and likely make significant strides in, among other fields, large aircraft development and hypersonic vehicle technology. Over the fifteen years that followed, the sequential Five-Year Plans for Science and Technology Development provided updates to the PRC’s strategic-level goals. More tactically, the leading organizations for the Major Projects reflected the institutional structures of the relevant PRC S&T ecosystems – while implementation plans they published, lines of effort they funded, and success cases they identified illustrated more specific technical objectives, investments, and capabilities. The same remains true for the current generation of Major Projects.



Of course, the PRC does obfuscate some of the relevant information. For example, the hypersonic vehicle technology Major Project was not explicitly discussed in any public-facing PRC government plans. But details about it do exist in the open source. And other, military-relevant elements of the PRC planning system are not obfuscated the same way. By virtue of the PRC's national-level strategy of military-civil fusion, most of the country's scientific, technological, and industrial initiatives, including those that are openly discussed and are not classified, are military-relevant. The Large Aircraft Major Project was responsible for developing the Y-20 transport aircraft. The Core Electronic Devices, High-End General-Purpose Chips, and Basic Software Products Major Project supported the PKS System, a PRC domestic alternative to the "Windows+Intel" system that has been used across the PRC military apparatus (New Times Securities, 2021).

United States understanding and monitoring of the PRC's planning system – and the capability development and investments it drives – should interface directly with the defense acquisition process. For example, early warning of PRC ambitions should be used to shape investments into responsive capabilities, to vet supply chains supporting acquisition priorities for exposure to Chinese dominance and influence upstream; and, to guarantee that requirements definition and analysis-of-alternatives processes receive threat intelligence inputs as they set the direction for major defense acquisition programs. Identification of entities involved in operationalizing China's ambitions should be used to ensure sufficient barriers against PRC bad actors in the acquisition process and, more broadly, across Federal procurement. Nor should this be just a defensive dynamic. Monitoring of PRC capability development should be used to ascertain weaknesses that can be targeted with the defense acquisition process and the integrated priority lists of operational defense acquirers, including the Combatant Commands, should be measured according to such a logic.

## **Understanding the Planning System to Shape the PRC**

Understanding the PRC's planning system offers more than just insight into the adversary's goals, capabilities, and key entities. It also offers insight into the ways in which the PRC responds to external stimuli – including the types and timing of stimuli that generate responses, as well as the nature of those responses. With that insight, the United States has the potential not only to monitor but also, proactively, to shape the PRC system.

For all its institutionalization, the PRC's planning system is not rigid. Rather, it leaves room for flexibility and consistently responds to external stimuli. As already noted, both the slate of 2006 Major Projects and their timing were informed by international developments in the early 2000s that changed Beijing's perception of the geopolitical environment. Foreign countries (e.g., Israel) and companies (e.g., Microsoft) imposed restrictions on access to their technologies that exacerbated the PRC's concerns about international dependence, therefore amplified the country's emphasis on self-reliance, and in turn led it to accelerate its investments to that end.

The Innovation 2030 Major Projects, first established in 2016, saw an even more drastic systemic change, this one driven by external technological trends. The original set of Innovation 2030 Major Projects, listed in the 13th Five-Year Plan for Science and Technology Innovation (“十三五”国家科技创新规划), was composed of 15 publicly disclosed initiatives. But in 2017, the State Council added a sixteenth: the New Generation AI Major Project (aka Artificial Intelligence 2.0 Major Project). This late-addition Major Project appears formally to have been established via the New Generation Artificial Intelligence Development Plan (新一代人工智能发展规划). And that plan constituted a response to emerging PRC diagnoses regarding artificial intelligence as a transformative technology – informed at least in part by Chinese Academy of Engineering research and pressure during the 2015 to 2017 period (China Net, 2017).





These cases indicate that the PRC planning system responds to perceived threats (e.g., technology restriction) and opportunities (e.g., emerging technologies with potentially disruptive capabilities). These cases also indicate that the timing of those stimuli may inform the PRC's response. At the most basic level, opportunities or threats identified in planning's research process or before that process begins – like concerns over Israel's provision of aerospace technology – can readily be incorporated into the process's outputs (e.g., converted into Major Projects). Opportunities or threats identified after the planning process ends, like Microsoft's "black screen," may be likely only to generate marginal changes (e.g., speed the implementation of relevant Major Projects). And opportunities or threats identified before plan finalization but after research phases (e.g., the emergence of AI as a disruptive technology) may be able to be integrated into the research pipeline and subsequent resourcing, but with a lag and therefore at greater cost. For example, CAE was able to initiate research on AI's transformative potential before the 13th Five-Year Plan was formalized in 2016. But the results were not incorporated into the Major Project system until after that plan was published.

Understanding these factors – the stimuli that the PRC's planning system responds to, how they interface with different steps of the process, and how their effects might vary at different points in the process – can all support development and deployment of a proactive, adversary informed defense acquisition process. Specific interventions from the United States, and its defense acquisition process, may be particularly efficacious at specific points in the Chinese planning process.

The process of applying this framing to defense acquisition could begin with a rubric that maps defensive and offensive actions onto different points in the Chinese planning system, namely "early," pre-capability points, when the Chinese planning process has formally identified areas in which to invest but has not yet scaled resourcing of them, and "late," post-capability points, when investments have been completed to field a new capability.

Defensive actions early in the PRC planning process, when new Chinese investment areas have been formalized and are being scaled, could focus on updating existing US capabilities proactively to defend against emerging PRC priorities. Doing so would allow the United States to be prepared when the PRC ultimately demonstrates its capabilities, and by extension to be more efficient in its investments and corresponding operational capability development. For example, if Beijing were to launch a program to develop a new anti-ship ballistic missile that would increase the PRC's ability to restrict US maneuver in the western Pacific, the United States might respond by leveraging a mix of procurement and operations and maintenance (O&M) funding and modifying existing programs of record to arm surface vessels with improved camouflage and concealment technology and capabilities.



|       | Defense                                                               | Offense                                                     |
|-------|-----------------------------------------------------------------------|-------------------------------------------------------------|
| Early | Use O&M funding to update existing capabilities                       | Use R&D funding to put new PRC investments on the defensive |
| Late  | Re-assert escalation dominance by substituting operational approaches | Demonstrate disruptive offensive capabilities               |

Offensive actions early in the PRC planning process could focus on R&D investment that forces the PRC to defend its new investments. For example, were Beijing to launch new programs focused on drone swarms, leading innovation forces within the Department of War, like DARPA and the Defense Innovation Unit, could launch programs to research and develop systems able offensively to target the networking hubs and coordinating mechanism required for those drone swarms. Actions like this – and signals of them – would force the PRC to expend additional resources defending its drone swarms, while also bolstering US prospective responses to them.

Defensive actions late in the PRC planning process, when capabilities have been developed, could orient around re-asserting escalation dominance by changing operational approaches. Take, for instance, PRC investments into controlling supply of a key resource, like rare earth elements. Once that capability is developed, the US might defensively respond by revealing the ability to substitute for that resource (e.g., motors that do not require rare earth permanent magnets). The substitute in its own right would help to protect the United States from the new PRC threat. By demonstrating it late, the United States would have encouraged the PRC to complete its investments in its now-less-threatening capability, and therefore to expend additional resources.

Finally, offensive actions late in the PRC planning process could feature demonstration of disruptive offensive capabilities that invalidate Chinese investments. In the same supply chain case, for example, the United States might unveil an electromagnetic pulse weapon able to target counter-value industrial base capabilities across the Chinese mainland – therefore placing at risk the basic premise, and foundation, of the new PRC capability.

This is a notional and generalized rubric. But, at a minimum, it demonstrates a theoretical model according to which an adversary-informed, proactive defense acquisition system could orient – offensively and defensively – to target the PRC planning system. At the same time, this abstracted rubric also underscores the need to incorporate threat intelligence earlier in, and in a recurring fashion throughout, the US capability development process.



## Works Cited

- Harper, Jon. (March 7, 2022). "U.S. Challenged to Defend Against Chinese Missiles". National Defense Magazine.
- Yao, Yichuan. (November 27, 2012). "The Measures Clarify the 'Two Categories and Four Levels' Planning System and the Relationship between Various Types and Levels of Planning [《办法》明确了'两类四规划体系及各类各级规划之间的关系 中华人 共和]". Ministry of Natural Resources of the People's Republic of China.
- Yang, Ly. (2011). "Research on the Organization Methods and Effects of Group Decision Making in National Special Planning [国家级专项规划群体决策的组织方式及其果研究]". Public Administration.
- Yu, Dong. (April 28, 2025). "The History, Experience, and Prospects of a Unified Planning System [统一规划体系的历程、经验和展望]". Guangming Daily.
- "微软黑屏：国产软件的机遇？ [Microsoft's 'Black Screen:' An Opportunity for Domestic Software?]." (November 6, 2008). Science Times.
- "Notice of the General Office of the State Council on Adjusting the Composition of the National Science and Technology Education Leading Group [国务院办公厅关于调国家科技教育领导小组组成人员的通知]". (May 2008). State Council of the PRC.
- "世界第二 中国突破高超声速冲压发动机自主飞行 [World's Second: China Achieves Autonomous Flight Breakthrough with Hypersonic Scramjet Engine]". (October 9, 2015). Paris Labs.
- "深度解读：《新一代人工智能发展规划》（附原文） [In-Depth Analysis: 'The Next Generation Artificial Intelligence Development Plan' (Original Text Included)]." (September 2, 2017). China Net.
- "'东风-17'：高超声速让反导系统形同虚设 ['Dongfeng-17:' Hypersonics Render Missile Defense Systems Useless]". (January 15, 2018). Science and Technology Daily.
- "中国商飞：大飞机翱翔蓝天惊艳世界 看中国浮沉'民机梦' [COMAC: Large Aircraft Soaring Through the Blue Sky, Amazing the World: Witnessing China's Rise and Fall in the 'Civil Aircraft Dream']". (October 18, 2018). Shanghai Municipal Commission of Economy and Informatization.
- "'PK'体系进展超预期·2021 年政策驱动明显 [The Progress of the 'PK' System Exceeded Expectations, and the Policy Drive in 2021 Was Obvious]". (January 4, 2021). New Times Securities.
- "China Surprises U.S. with Hypersonic Missile Test, FT Reports." (October 17, 2023). Reuters.
- "一小时打遍全球！西工大'飞天一号'成功试射，超级强国的利器 [Strike Anywhere on Earth Within an Hour! NWPU's 'Feitian-1' Successfully Test-Fired—A Superpower's Ultimate Weapon]". (July 6, 2022). Guancha.
- "This Is How China's Five-Year Plans Are Formulated and Implemented, [中国五年规划是这编制实施的]". (August 1, 2025). China Daily.
- "卫生部：'重大新药创制'科技重大专项正式启动 [Ministry of Health: 'Major New Drug Innovation' National Science and Technology Major Project Officially Launched]". (March 17, 2026). Unitalen Attorneys at Law.
- "'东风-17'：高超声速让反导系统形同虚设 ['Dongfeng-17:' Hypersonics Render Missile Defense Systems Useless]". (January 15, 2018). Science and Technology Daily.



# Defense Acquisition and Contracting Approaches: Implications for Foreign Military Sales

**John G. (Jerry) McGinn, PhD**—is the director of the Center for the Industrial Base at CSIS and a senior fellow in defense and security. A leading expert on industrial base and acquisition policy, he previously led the Baroni Center for Government Contracting at George Mason University, producing over 70 reports. He served in the Department of Defense analyzing the defense industrial base and overseeing merger reviews and CFIUS transactions. His career includes roles at Northrop Grumman, Deloitte, QinetiQ, and RAND. A former Army infantry officer, he has testified before Congress and holds degrees from Georgetown and West Point.

**Celia Barrie**—is a program coordinator and research assistant for the Center for the Industrial Base at the Center for Strategic and International Studies (CSIS). Her analytical focus includes data management, database architecture, and quantitative research, with experience supporting analysis through satellite imagery, OSINT collection, and geospatial information systems. Prior to joining CSIS, she worked as a geospatial data intern supporting mapping projects and conducted GIS-based research through the UCLA Geography Department. She holds a BA in Geography and a BA in Global Studies, with a minor in Geospatial Information Systems, from the University of California, Los Angeles.

## Abstract

This paper examines how contracting structures shape the United States' Foreign Military Sales (FMS) system. It challenges explanations that attribute FMS inefficiencies to the limited use of acquisition "innovation pathways," such as Other Transaction Authority (OTA) or the Middle Tier of Acquisition (MTA), and instead argues that outcomes are better explained by established Federal Acquisition Regulation (FAR)-based contracting instruments. Using data from the Federal Procurement Data System (FPDS), the analysis compares FMS and non-FMS procurement across contract types, platform portfolios, vendor composition, and time trends.

The findings show that FMS contracting is concentrated in mature production and sustainment activities, dominated by Firm-Fixed-Price (FFP) contracts with limited but growing use of incentive-based structures. Indefinite Delivery/Indefinite Quantity (IDIQ) contracts play a supporting role, while Undefined Contract Actions (UCAs) remain marginal. Compared to domestic procurement, FMS relies more heavily on large contractors, reflecting both the scale and complexity of procured systems as well as institutional constraints such as pricing transparency and intergovernmental coordination.

Overall, the study reframes FMS as a constrained system in which performance is shaped less by new acquisition authorities and more by how conventional contracting tools are applied, highlighting opportunities for reform through improved implementation of existing mechanisms.

## Introduction

The United States' Foreign Military Sales (FMS) system sits at the core of modern security cooperation, serving as a primary mechanism for the United States to equip its allies, strengthen partnerships, and advance its strategic interests abroad. Managed by the Defense Security Cooperation Agency (DSCA), FMS enables partner nations to acquire U.S. defense capabilities while promoting interoperability and long-term institutional alignment (DSCA, n.d.). As great power competition intensifies and technological change accelerates, the effectiveness of this system has taken on renewed importance (Byman et al., 2025). Despite its strategic value, however, FMS continues to face persistent criticism for long delivery timelines, bureaucratic complexity, and limited flexibility in responding to evolving partner requirements (Saum-Manning et al., 2024).

These challenges are not unique to FMS, but instead reflect broader structural features of the U.S. defense acquisition system. Over time, acquisition reform efforts have sought to



improve speed, predictability, and efficiency in how the Department of Defense procures and sustains capabilities (Lopez, 2023). While these efforts have introduced more flexible contracting practices in certain domestic acquisition contexts, their translation into FMS has been uneven. Rather than reflecting a failure to adopt “innovative pathways” such as Middle Tier of Acquisition or Other Transaction Authority, this gap is better explained by the nature of FMS itself. FMS primarily executes the sale of fully-developed systems, where procurement is governed less by research and development and more by established contracting fundamentals including cost certainty, compliance, and delivery schedule.

Existing research has documented persistent inefficiencies in FMS, including longer delivery timelines, reduced responsiveness, and coordination challenges across multiple stakeholders. However, much of the policy discourse remains high-level, emphasizing broad reform goals without systematically analyzing the distribution and performance of specific contracting mechanisms. There is limited empirical attention to how core contract types—such as Firm-Fixed-Price contracts, incentive-based structures, and Indefinite Delivery/Indefinite Quantity (IDIQ) vehicles—are used across FMS compared to non-FMS procurement and how these choices shape execution outcomes.

Ultimately, improving the FMS process is not primarily a matter of adopting new acquisition “innovation tools,” but of better understanding and optimizing the contracting structures already in use. As global demand for defense capabilities increases, the ability of the United States to deliver timely, predictable, and scalable solutions to partners will depend on how effectively these contracting mechanisms are employed. This research therefore links acquisition practice to security cooperation outcomes by providing an evidence-based assessment of how contract structure shapes performance in FMS relative to broader defense procurement.

## Approach

This paper makes three core contributions to the study of Foreign Military Sales (FMS) and defense acquisition reform. First, it reframes FMS performance not as a function of acquisition pathway innovation, but as an outcome of how established Federal Acquisition Regulation (FAR)-based contracting instruments are selected and applied. Second, using Federal Procurement Data System (FPDS) analysis, it demonstrates that FMS contracting is structurally concentrated in mature production and sustainment activities governed primarily by Firm-Fixed-Price and related contract types (Warfighting Acquisition University, n.d.). Third, it shows that variation in FMS outcomes is better explained by contract structure, vendor composition, and system maturity than by the adoption of alternative acquisition authorities such as OTA or MTA. Together, these findings shift the analytical focus of FMS reform from procedural innovation toward optimization of existing contracting mechanisms.

To support this analysis, the study employs a mixed-methods approach. Quantitatively, it draws on data from the Federal Procurement Data System (FPDS) to examine millions of defense contracts, including those associated with FMS cases. This enables a comparison of contract type distribution, platform portfolios, contract size, and duration across FMS and non-FMS procurement. Qualitatively, the study incorporates case-based vignettes to contextualize observed contracting patterns and assess how different mechanisms function in practice within established defense systems. A central contribution of this research is the development of a data-driven framework that allows practitioners within the security cooperation workforce to better understand historical contracting patterns and apply those insights to current FMS cases. By focusing on widely used contracting instruments rather than R&D-oriented acquisition pathways, the analysis provides more directly applicable insights for decision-makers operating within existing institutional constraints.





## Explaining the Shift in Project Focus

An initial review of Federal Procurement Data System (FPDS) records, supplemented by DSCA case-level information, suggests that Foreign Military Sales (FMS) contracting is focused on the production, delivery, and sustainment of established defense systems. Across major platform portfolios, including aircraft, missile systems, and land vehicles, FMS activity is primarily associated with procurement decisions that occur after system development has already been completed within the domestic acquisition system. In other words, FMS is not typically a vehicle for early-stage acquisition activity (Warfighting Acquisition University, n.d.). Instead, it functions as a downstream execution mechanism for capabilities that are already production-ready and integrated into existing U.S. defense supply chains (Warfighting Acquisition University, n.d.). Because FMS and domestic acquisition rely on the same underlying Federal Acquisition Regulation (FAR)-based contracting tools, differences in outcomes can be traced to how those tools are designed, selected, and applied across different institutional environments. As a result, the underlying contracting activity reflects recurring production runs, sustainment support, spares procurement, and configuration adjustments rather than innovation or developmental work.

This empirical structure is important because it constrains the range of contracting behaviors observable in the data. The FPDS record shows limited evidence of FMS participation in prototyping or iterative capability experimentation. These categories of acquisition activity are more commonly associated with domestic U.S. procurement prior to foreign release decisions, rather than with the FMS process itself. This empirical reality has direct implications for how FMS should be analyzed in the context of acquisition reform and contracting behavior.

Prior reform efforts in the FMS space have consistently identified inefficiencies related to delays, coordination complexity, and limited transparency. For example, structural analyses of the security cooperation process highlight bottlenecks in pricing and approval workflows that slow execution, while government-led modernization efforts have emphasized the need to streamline interagency coordination and reduce administrative friction across the enterprise (Sanders et al., 2019; U.S. GAO, 2017). Similarly, defense industry assessments point to fragmented oversight and inconsistent execution standards as persistent challenges, and broader policy discourse reinforces the need for improved governance and process efficiency across security cooperation activities (Saum-Manning et al., 2024). While these perspectives provide valuable insights into procedural and institutional constraints, they remain largely focused on administrative reform and do not systematically account for how contracting structure itself shapes execution outcomes in FMS. This limitation motivates a shift toward analyzing the distribution and use of FAR-based contracting instruments as the primary explanatory lens, rather than the initially proposed assessment of novel acquisition pathways.

First, acquisition pathways such as the Middle Tier of Acquisition (MTA) are designed primarily to support rapid prototyping and accelerated fielding of new capabilities (Warfighting Acquisition University, n.d.). These pathways assume a level of technical uncertainty and iterative development that is generally not present in FMS cases, where the underlying system has already been developed, tested, and placed into production (U.S. GAO, 2023). As a result, MTA is structurally misaligned with the dominant characteristics of FMS procurement activity and does not serve as a meaningful explanatory variable for differences in FMS contracting outcomes.

Second, Other Transaction Authority (OTA) mechanisms are similarly limited in their relevance to FMS analysis. OTAs are primarily used to facilitate prototyping, non-traditional vendor engagement, and flexible experimentation outside the constraints of the Federal Acquisition Regulation (FAR; Office of the Under Secretary of Defense for Acquisition and



Sustainment, 2023). While these tools have become increasingly important in certain areas of domestic defense innovation, they are not representative of the contracting environment that defines FMS execution. FMS cases must comply with stricter requirements for transparency, pricing documentation, intergovernmental coordination, and congressional oversight, all of which significantly constrain the applicability of OTA-style contracting structures (U.S. GAO, 2025).

These considerations suggest that both MTA and OTA are analytically peripheral to the core dynamics of FMS contracting behavior. While they are important components of the broader acquisition reform ecosystem, they do not meaningfully capture the structure of how FMS cases are executed in practice. Because FMS is centered on the execution of demonstrated defense systems, acquisition pathways designed for prototyping and experimentation are not strong explanatory variables for variation in FMS outcomes. Their limited presence in FMS-related contracting activity reduces their utility for systematic analysis of timelines, contract structures, or performance differences between FMS and non-FMS procurement.

*Figure 1* illustrates the distribution of total obligated funding across products, research and development (R&D), and services for both FMS and non-FMS contracts over time. The data shows that FMS obligations are saturated in product procurement, with only a negligible share allocated to R&D activities. In contrast, non-FMS contracting includes a more visible, though still limited, R&D component. This disparity highlights a fundamental structural difference: FMS operates primarily as a mechanism for acquiring and delivering well-developed systems rather than developing new capabilities. As a result, contracting approaches designed for prototyping and experimentation play a minimal role in FMS execution.




**Figure 21. Total Amount Obligated by Product, Service, and R&D**

Accordingly, the analytical focus of this study shifts away from innovation-oriented acquisition pathways and toward the contracting structures that actually govern FMS execution.

This includes widely used and empirically observable mechanisms such as Firm-Fixed-Price contracts, incentive-based structures, and Indefinite Delivery/Indefinite Quantity (IDIQ) vehicles, as well as the institutional constraints that shape their use. By refocusing on these established contracting instruments, the analysis is better aligned with the empirical reality of FMS and more capable of generating actionable insights for improving execution within existing institutional boundaries.

## **What Actually Drives FMS Contracting Outcomes?**

Rather than being shaped by experimental acquisition pathways, FMS contracting outcomes are primarily driven by a set of well-established FAR-based contracting instruments. These mechanisms structure how work is priced, how risk is allocated between the U.S. government and contractors, and how execution is managed over a contract's lifetime. To assess how these FAR-based contracting mechanisms shape FMS efficiency in practice, this study examines several widely used contracting instruments.

The most common instrument, utilized by both FMS and non-FMS cases is the Firm-Fixed-Price (FFP) contract. In the FMS context, FFP contracts are frequently used from mature systems such as aircraft, missiles, and vehicles, where production costs are relatively stable and requirements are well understood (GSA, 2026). FFP contracts place most cost risk on the contractor while providing the government and, by extension, foreign partners with budget certainty (GSA, 2026). In contrast, cost-based contracts, such as Cost Plus Award Fee, are used in situations where requirements are less defined or where cost uncertainty remains high (GSA, 2026). Under these arrangements, the government agrees to reimburse allowable costs incurred by the contractor, often with additional fees (GSA, 2026). While less common in stable production environments, cost-type contracts may still appear in early production lots, complex sustainment efforts, or cases involving technical uncertainty (GSA, 2026).

A third category this study investigates includes incentive-based contracts, such as Fixed-Price Incentive Firm (FPIF) and Cost-Plus Incentive Fee (CPIF) structures. These contracts are designed to manage execution risk by linking contractor compensation to performance outcomes such as cost control, schedule adherence, or technical achievement (GSA, 2026). Their use reflects a constrained form of flexibility: While they introduce performance incentives, they remain embedded within a broader institutional environment that prioritizes price certainty, transparency, and intergovernmental accountability (GSA, 2026). As a result, incentive-based contracting represents incremental adjustment within a fundamentally structured system rather than a shift toward experimental acquisition approaches.

Another widely used mechanism is the Indefinite Delivery/Indefinite Quantity (IDIQ) contract structure. IDIQs enable the government to issue task or delivery orders over time without re-competing or renegotiating the underlying contract, providing a degree of administrative and logistical flexibility within the FAR framework (GSA, 2026). This structure is particularly well suited to sustainment, spare parts, and recurring support requirements where demand is continuous but not fully defined at the outset (GSA, 2026). IDIQs are not unique to FMS and instead reflect a standardized contracting instrument across the defense acquisition system. In Fiscal Year 2024, over half of all Department of Defense obligated dollars were on Indefinite Delivery Vehicles (IDVs), which include IDIQs (Congressional Research Service, 2024). Their widespread use suggests that flexibility in both FMS and non-FMS contexts is achieved primarily through established FAR-based tools rather than through acquisition pathway variation or specialized procurement authorities.

Finally, Unfixed Contract Actions (UCAs) represent a more exceptional contracting instrument that permits performance to begin before final contract terms are fully agreed upon (GSA, 2025). While UCAs can be used to accelerate execution in urgent operational contexts,



their application in FMS is structurally constrained by requirements for pricing justification, intergovernmental coordination, and eventual contract definitization (GSA, 2025). More broadly, the limited role of UCAs across both FMS and non-FMS contracting underscores their function as an exception-handling mechanism rather than a routine feature of acquisition strategy. Their marginal use reinforces the broader finding that execution speed and flexibility in FMS are not driven by ad hoc contracting authorities, but are instead shaped by the disciplined application of established contracting structures within a highly regulated environment.

These contracting mechanisms are more analytically relevant than acquisition “innovation pathways” such as the Middle Tier of Acquisition (MTA) or Other Transaction Authority (OTA) because they directly shape the structure and performance of FMS execution. Contract type governs four core dimensions of acquisition outcomes: price certainty, schedule predictability, risk allocation, and execution flexibility (GSA, 2026). Fixed-price contracts provide cost stability for both the U.S. government and foreign partners, while contract structure more broadly influences delivery timelines, the distribution of financial risk, and the degree of flexibility available during execution (GSA, n.d.). Incentive contracts balance risk between stakeholders, while mechanisms such as IDIQs and UCAs introduce varying levels of adaptability (GSA, n.d.).

These factors shape FMS performance far more directly than whether a program falls under a particular acquisition pathway. This leads to a fundamental reframing of the research question. Rather than asking which acquisition pathway is used, the more relevant question is: Which contract structure governs the execution of an already-developed system within FMS? By shifting the focus from acquisition pathways to contracting structures, this analysis better reflects the operational reality of FMS and provides a clearer understanding of what drives differences in cost, schedule, and performance outcomes. Ultimately, FMS performance is shaped less by innovation in acquisition design and more by the selection and execution of conventional contracting tools within a highly structured environment.

## **FPDS Analysis of FMS by Contract Type**

This section presents a series of static visualizations derived from the Federal Procurement Data System (FPDS) data comparing FMS and non-FMS contracting activity across multiple dimensions. The accompanying analysis provides an initial empirical mapping of how contracting structures differ across these two procurement environments. Rather than focusing on acquisition pathways or experimental contracting mechanisms, the figures examine observable variation in contract type distribution, platform portfolios, vendor size, and time trends in obligated funding.






**Figure 22: Total Amount Obligated by Contracting Method**

These cuts establish a descriptive baseline for understanding how FAR-based contracting instruments are applied in practice across FMS and non-FMS contexts. The contract type distribution provides the clearest evidence of structural differences in contracting behavior. *Figure 2* demonstrates the distribution of FAR-based contracts over the last decade. Across the dataset, the majority of FMS obligations are Firm-Fixed-Price contracts; however, there is evidence of gradual diversification, as the use of incentive-based contract structures shows modest growth. The large proportion of FFP contracts suggests a preference for cost certainty, clearly defined deliverables, and reduced pricing ambiguity in foreign partner procurement. However, the growth in incentive-based contracts indicates an increasing effort to introduce more flexible arrangements that can better balance cost control with performance incentives in cases of greater execution uncertainty. When examining non-FMS cases, the distribution of contract types appears more balanced across categories, with gradual increases in Firm-Fixed-Price contracts as well as modest shifts in other contract types. Proportionally, this suggests a relatively stable but less specialized contracting structure compared to FMS, with greater dispersion across contracting methods reflecting a broader range of procurement requirements and use cases.

These findings have important implications for FMS reform. The continued concentration of Firm-Fixed-Price contracts suggests that FMS contracting is structurally oriented toward cost certainty, standardization, and the transfer of execution risk to contractors, reflecting the requirements of intergovernmental coordination and pricing accountability. While the modest increase in incentive-based contracting indicates a gradual move toward greater flexibility, this shift remains incremental rather than transformative and does not meaningfully alter the overall dominance of fixed-price arrangements. As a result, reform efforts are unlikely to be driven by wholesale changes in contract type usage, but rather by improvements in how existing FFP-heavy structures are priced, negotiated, and executed. In contrast, the more diversified contract mix observed in non-FMS procurement reflects a broader range of mission requirements and greater heterogeneity in acquisition needs, rather than a clearly superior model of contracting



flexibility. Accordingly, the most meaningful gains in FMS performance are likely to come from optimizing execution within the constraints of the existing contracting structure rather than attempting to replicate the distributional patterns of non-FMS procurement.

Additional cuts by platform portfolio further contextualize these differences by linking contract structures to specific capability areas, including aircraft, land vehicles, and ordnance and missiles. This dimension allows for assessment of whether contracting patterns are driven primarily by system type or by the institutional context of FMS versus non-FMS execution. Macro-level evidence indicates that mature systems are consistently associated with more standardized contracting approaches.

*Figure 3* reinforces this interpretation by disaggregating contract type distribution across platform portfolios in both FMS and non-FMS contexts. Across all categories, Firm-Fixed-Price contracts dominate, underscoring the overall stability and standardization of defense procurement. In ordnance and missile systems, FMS contracting is highly concentrated in FFP arrangements, with only a modest presence of incentive-based contracts, while non-FMS cases remain similarly FFP-heavy but exhibit a slightly more uniform distribution. Aircraft procurement shows the most notable variation, where FMS cases include a visible share of Incentive-Based contracts alongside dominant FFP structures, whereas non-FMS aircraft contracting remains largely FFP-driven with a more tightly clustered distribution overall. Land vehicles display the most consistent pattern across both channels, with overwhelming reliance on FFP contracts and minimal variation in contract type. Collectively, these patterns indicate that while system complexity introduces limited variation at the margins, FMS contracting remains fundamentally standardized across platform portfolios and is primarily defined by stable, production-oriented contracting norms.




**Figure 3: Total Obligated Amount by Contracting Method for Various Platforms**

Vendor size analysis provides a complementary perspective on how industrial base structure interacts with contracting behavior. *Figure 4* demonstrates the pattern of contracting methods within various vendor sizes. By examining the distribution of obligations across large

prime contractors and smaller vendors, the analysis identifies the extent to which contracting activity is focused among major defense firms. Initial results suggest that FMS contracting remains highly reliant on established primes, consistent with the scale, integration requirements, and compliance demands of international defense procurement. Non-FMS activity, while still dominated by large contractors, shows comparatively greater dispersion across vendor categories, reflecting a broader range of sourcing strategies within domestic acquisition.




**Figure 4: Total Amount Obligated by Contracting Method for Various Vendor Sizes**

This distributional pattern varies meaningfully once disaggregated by contract type. Among large contractors, including the “Big 5” and other major primes, FMS activity is characterized by a mix of Firm-Fixed-Price (FFP) and incentive-based structures, with a relatively stronger presence of Incentive-Based arrangements in FMS relative to non-FMS cases. This suggests that large-scale foreign military procurement incorporates performance-aligned contracting mechanisms where program complexity or production scale warrants additional flexibility, even within a system otherwise oriented toward cost certainty. In contrast, non-FMS contracting with large vendors remains proportionally more centered in FFP arrangements, though still displaying diversity across contract types, including a non-trivial share of cost-based structures.

A different pattern emerges in the medium and small vendor categories. FMS contracting at the medium-tier level is still largely dominated by FFP structures but includes a small presence of cost-based contracts, indicating selective flexibility in areas where requirements uncertainty or sustainment complexity is higher. Non-FMS contracting with medium and small vendors is heavily skewed toward FFP instruments, with comparatively limited variation in contract type selection. Small vendor activity is concentrated in FFP arrangements across non-FMS procurement, reinforcing the tendency for standardized pricing and reduced contractual complexity at lower tiers of the industrial base.

Time-series trends offer additional insight into how these structural patterns evolve over the FY 2010–FY 2024 period. These visualizations show that while FMS contracting remains

consistently dominated by FFP structures, there is modest and gradual diversification over time, particularly through incremental increases in incentive-based contracting and stable but limited use of IDIQ mechanisms. Non-FMS contracting similarly exhibits persistence in its underlying structure, with FFP remaining central but accompanied by slightly broader variation across contract types. Importantly, IDIQ usage remains consistently small in both environments, suggesting that while it provides flexibility in execution, it does not constitute a primary driver of contracting composition. Overall, the stability of these distributions reinforces the conclusion that neither FMS nor non-FMS procurement is characterized by frequent structural shifts in contracting approach over time; rather, differences are persistent and embedded in institutional and programmatic conditions.

Figure 5 illustrates the prevalence of Indefinite Delivery Contracts (IDCs) within both FMS and non-FMS procurement. Across both FMS and non-FMS environments, Indefinite Delivery/Indefinite Quantity (IDIQ) contracts make up a significant portion of contracting activity and function as flexible mechanisms for managing recurring procurement and sustainment requirements. This pattern indicates that flexible, indefinite delivery structures are not peripheral tools, but rather a significant component of how procurement is executed in practice. While IDCs can include IDIQ arrangements, they are not perfectly synonymous due to FPDS classification limitations. Nevertheless, the consistent presence of IDC-designated contracts across both FMS and non-FMS contexts demonstrates that flexible contracting mechanisms are already widely implemented within the broader acquisition system. Importantly, these vehicles are not unique to FMS, but their established use elsewhere suggests they could be more deliberately leveraged to address variability in demand and improve execution efficiency within the FMS framework.




**Figure 5: Total Amount Obligated by Indefinite Delivery Contract (IDC)**

Figure 6 further contextualizes the role of Unpriced Contract Actions (UCAs) by examining their share of obligated funding over time in both FMS and non-FMS environments. The data shows that UCAs consistently represent a small portion of total obligations in both contexts, even as overall funding levels fluctuate. In FMS, UCA usage is slightly more visible, particularly in years with elevated obligation levels, reinforcing their role as an exception-driven mechanism used primarily in situations where “the government needs the contractor to start

work immediately and there is insufficient time to negotiate all terms” (U.S. GAO, 2007). However, their share remains limited relative to fully priced contracts, indicating that most FMS execution continues to rely on definitized agreements. In non-FMS procurement, UCAs are even less prominent, with only minimal contributions to total obligations across all years. These patterns reinforce the interpretation of UCAs as a narrowly applied, contingency-based instrument rather than a scalable mechanism for managing procurement, with limited influence on the overall structure of contracting activity in either environment.




**Figure 6: Total Amount Obligated by UCA Classification**

The analysis of IDIQs and UCAs shows that FMS contracting is primarily executed through established FAR-based instruments. IDIQ/IDC structures are consistently present in both FMS and non-FMS procurement, reflecting their role as standardized mechanisms for managing recurring requirements and sustainment activities within mature acquisition programs. UCAs remain a marginal instrument across both environments, appearing infrequently relative to fully definitized contract actions. Within FMS, UCAs occur more frequently than in non-FMS procurement but still represent a small share of overall obligations. This pattern indicates that UCAs function as a limited contingency mechanism rather than a routine contracting structure. Further FPDS analysis will examine more granular data to better characterize how IDIQs and UCAs are used across program contexts.

FPDS analysis conducted to-date demonstrates that observable differences between FMS and non-FMS procurement are primarily a function of how a stable set of FAR-based contracting instruments is deployed across distinct institutional and industrial contexts. Rather than revealing fundamentally different acquisition systems, the data shows variation in the distribution and application of a common contracting toolkit, particularly Firm-Fixed-Price, incentive-based, cost-type, and IDIQ structures. These instruments are consistently used to manage tradeoffs between cost certainty, execution flexibility, and risk allocation, with their relative mix shaped by system maturity, vendor composition, and the constraints of foreign partner coordination (GSA, 2026). Across all dimensions, contract type, platform portfolio, vendor size, and time trends, the FPDS data consistently point to structural persistence rather than event-driven variation.

These findings reinforce the broader argument that FMS contracting behavior is best understood through the interaction of contract structure and industrial base composition rather than through acquisition pathway variation or event-based contracting tools. Tools such as OTA and MTA play little to no structural role in shaping observed FMS contracting behavior, reflecting their misalignment with the execution-focused nature of foreign military sales. Similarly, UCAs appear only as a marginal instrument with limited impact on system-level variation. Instead, the primary drivers of divergence are found in the distribution of FAR-based contract structures, differences in vendor scale and industrial base composition, the maturity of underlying defense systems being procured, and the institutional constraints inherent in foreign procurement processes. These factors suggest that FMS is best understood not as an innovation-oriented acquisition pathway, but as a constrained optimization system in which established contracting instruments are applied within tighter informational, political, and compliance boundaries than in domestic procurement contexts.

## Conclusion

This paper has argued for a shift in how Foreign Military Sales (FMS) is analyzed, moving away from acquisition pathway narratives and toward a contracting structure framework. The analysis shows that differences between FMS and non-FMS procurement are primarily structured through the deployment of established Federal Acquisition Regulation (FAR)-based contracting instruments within a highly regulated institutional environment. Across contract type distributions, vendor structures, platform portfolios, and time trends, the evidence consistently indicates that contracting composition is the key observable dimension along which FMS and domestic procurement diverge.

The central finding of the analysis is that FMS contracting is organized around the constrained application of standard FAR instruments. Firm-Fixed-Price, incentive-based contracts, IDIQ vehicles, and related mechanisms collectively shape core dimensions of execution, including cost certainty, schedule predictability, and risk allocation. These instruments operate within institutional constraints specific to foreign procurement, including intergovernmental coordination requirements, pricing transparency expectations, and compliance obligations, which reinforce a system characterized by stability, standardization, and limited contractual flexibility.

These patterns suggest that variation in FMS outcomes is primarily explained by how FAR-based mechanisms are structured and operationalized within the FMS environment, rather than by differences in acquisition pathway design. As such, meaningful reform is most plausibly located within adjustments to the institutional conditions governing the use of existing FAR instruments—particularly how standard contract types are selected, configured, and executed in practice across FMS transactions.

Finally, this research remains descriptive in nature. The FPDS analysis provides a macro-level mapping of contracting structure across FMS and non-FMS procurement and does not extend to evaluating alternative acquisition frameworks or unobserved institutional mechanisms. The next phase of research will build on these findings through program-level case studies to further examine how these contracting structures operate in execution and how they relate to observed procurement outcomes.

## Next Steps

While the FPDS analysis provides a systematic, macro-level view of contracting patterns across FMS and non-FMS procurement, it remains fundamentally descriptive in nature. It identifies how FAR-based contracting instruments are distributed across contracting types, vendors, platforms, and time, but does not directly observe how these structures operate in





execution. To move beyond description, the next stage of the research focuses on validating whether these observed contracting patterns correspond to measurable differences in program-level outcomes.

The purpose of this next phase is to shift from aggregate statistical patterns to causal analysis. Specifically, the case study approach is designed to assess whether variation in FAR-based contract structures is associated with differences in delivery performance, pricing behavior, and execution flexibility. This transition allows the study to evaluate not only how contracting tools are distributed, but how they function in practice under real program conditions. The case selection emphasizes major defense platforms that are central to U.S. foreign military sales and established to capture production and sustainment dynamics rather than developmental experimentation. These include fighter aircraft programs such as the F-15, F-16, F/A-18, and F-35, various helicopter programs including UH-60, and missile systems, in particular AMRAAM and PAC-3. Together, these cases represent a range of high value, long-duration procurement and sustainment relationships within the FMS system.

Across these cases, the analysis will examine whether contract structure helps explain variation in key execution outcomes, including delivery stability, pricing behavior, modification frequency, and sustainment flexibility. Pulling qualitative data from DSCA Defense Trade and Arms Transfers FMS Program and associated sources, this study will not only identify isolated contractual events, but will understand how FAR-based contracting methods work in practice and shape performance over time. This provides a more granular view of how contracting decisions translate into operational outcomes within complex defense systems. The case studies serve as a mechanism-level validation of the FPDS findings. While the dataset established broad structural differences between FMS and non-FMS procurement, the case studies will examine whether these differences are reflected in actual program execution. In doing so, they bridge the gap between aggregate contracting patterns and real-world performance, reinforcing the central argument that FMS outcomes are primarily shaped by the application of established FAR-based contracting instruments within a constrained institutional environment.

## References

- Byman, D., Palmer, A., Aldisert, A., Carroll, H., Cook, C., Park, C., & Jones, S. (2025). *Improving U.S. cooperation with allies and partners*. Center for Strategic and International Studies.
- Congressional Research Service. (2024). *Indefinite delivery, indefinite quantity contracts*.
- Defense Security Cooperation Agency. (n.d.). Defense trade and arms transfers: Foreign military sales.
- GSA. (n.d.). Part 16 - Type of Contracts.
- GSA. (2025). DFARS 217.74, Uninitiated Contract Actions.
- GSA. (2026a). FAR 16.103, Negotiating Contract Type.
- GSA. (2026b). FAR 16.2, Fixed Price Contracts.
- GSA. (2026c). FAR 16.3, Cost-Reimbursement Contracts.  
<https://www.acquisition.gov/far/subpart-16.2>
- GSA. (2026d). FAR 16.4, Incentive Contracts.
- GSA. (2026e). FAR 16.5, Indefinite-Delivery Contracts.
- GSA. (2026f). FAR Part 1, Federal Acquisition Regulations System.



- Lopez, C. T. (2023). *Tiger team recommendations aim to optimize foreign military sales*. Department of War.
- Office of the Under Secretary of Defense for Acquisition and Sustainment. (2023). *Other transactions guide*.
- Sanders, G., Akca, A., Halterman, N., Graham, J., & Cohen, S. (2019). *The effects of exporting on defense acquisition outcomes: A quantitative look at FMS contracting*. Acquisition Research Program.
- Saum-Manning, L., Marquis, J., Chindea, I., Elinoff, D., Pigott, T., & Brennan, E. (2024). *Optimizing foreign military sales roles*. RAND.
- U.S. GAO. (2007). *Defense contracting: Use of undefinitized contract actions understated and definitization time frames often not met*.
- U.S. GAO. (2017). *Foreign military sales: DOD needs to improve its use of performance information to manage the program*.
- U.S. GAO. (2023). *Middle-tier defense acquisitions: Rapid prototyping and fielding requires changes to oversight and development approaches*.
- U.S. GAO. (2025). *Other transaction agreements: Improved contracting data would help DoD assess effectiveness*.
- Warfighting Acquisition University. (n.d.-a). *International acquisition - Foreign military sales (FMS)*.
- Warfighting Acquisition University. (n.d.-b). *Middle tier of acquisition (MTA): MTA overview & benefits*.



## PANEL 9. FUTURE-PROOFING THE LIFECYCLE: SILO-BUSTING THROUGH INTEGRATED DIGITAL MATURITY

Wednesday, May 6, 2026, 1445 – 1600 ET (1145 - 1300 PT)

### Panel Summary:

Traditional, linear acquisition models struggle to keep pace with the volatile, interconnected nature of modern warfare. This panel introduces transformative frameworks designed to dismantle organizational silos and accelerate the delivery of mission-critical capabilities. From the "Quantum Acquisition" paradigm that reveals hidden interdependencies across complex architectures to AI-augmented "Human-in-the-Loop" systems that function as force multipliers for contracting officers, these research initiatives provide the tools necessary to operationalize the 2025 Acquisition Transformation Strategy. By blending advanced maturity modeling with real-time AI feedback, this panel offers a scalable roadmap for reducing procurement lead times while strengthening the strategic judgment of the acquisition workforce.

**Chair: Rear Admiral Kurk Rothenhaus, USN**, Commander, Naval Information Warfare Systems Command

### Presenters:

**Quantum Acquisition: A New Paradigm for Understanding the Interdependencies of Complex Networks in the Acquisition Life Cycle** – *Raymond Jones, Naval Postgraduate School*

**AI-Infused Integrated Digital Maturity Pathway (IDMP) User-Story Framework for Defense Acquisition Transformation** – *Darryl Draper-Amason, PhD, Old Dominion University*

**Augmented Acquisition: Operationalizing Human-in-the-Loop AI to Accelerate Warfighting Capabilities** – *Omar Haroun, Eudia*



**Rear Admiral Kurk Rothenhaus, USN**—is a native of New York City, New York. He received his commission in 1992 upon graduating from the University of South Carolina where he earned a Bachelor of Science degree. He also earned a Master of Science in Computer Science and a Ph.D. in Software Engineering from the Naval Postgraduate School, transferring to the Engineering Duty Officer community in 2003.

He assumed command of the Naval Information Warfare Systems Command (NAVWAR) in October of 2025. His previous flag assignments include Chief of Naval Research at the Office of Naval Research from June 2023, and Program Executive Officer, Command Control Computers Communications and Intelligence (PEO C4I) from May 2020.

He served as the Combat Systems/C5I officer on USS Harry S. Truman (CVN 75) and Chief Engineer on USS O'Brien (DD 975). Additionally, he served on the staff of Destroyer Squadron 15 embarked USS Kitty Hawk and on USS Fife (DD 991). He completed an Individual Augmentee tour in Baghdad, Iraq.

He completed numerous acquisition tours, including Program Manager for the Navy's Tactical Networks Program Office (PMW-160) and Commanding Officer of Space and Naval Warfare Systems Center Pacific. He also served as the Deputy Program Manager for the Navy Communications and GPS Program Office (PMW/A 170). Additionally, he served in various Assistant Program Manager roles in Command and Control, Battlespace Awareness, and Carrier Aviation Integration.



# **Quantum Acquisition: A New Paradigm for Understanding the Interdependencies of Complex Networks in the Acquisition Life Cycle of Warfighting Systems**

**Colonel (Ret.) Raymond Jones**—is the Professor of Practice at the Naval Postgraduate School. He served in the U.S. Army as a combat assault and experimental test pilot prior to entering the acquisition community. As an acquisition officer, he served as the Program Manager for two major defense acquisition programs and retired as the Deputy Program Executive Officer for the Joint Tactical Radio System. [rdjone1@nps.edu]

## **Abstract**

Defense acquisition programs continue to struggle with persistent cost overruns, schedule delays, and capabilities that fail to align with rapidly evolving threats. Traditional “Newtonian” linear models treat requirements, design, budgeting, and operations as semi-isolated domains, failing to capture the non-deterministic and highly interdependent nature of modern systems development. This paper introduces Quantum Acquisition, a metaphorical framework grounded in quantum mechanics, graph theory, and Bayesian inference. It models the defense acquisition life cycle as an entangled network of four architectural layers: system architecture, business/acquisition process architecture, Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, and Facilities (DOTMLPF), and the operational threat network. By viewing programs as probabilistic “acquisition probability clouds” rather than fixed paths, the framework highlights superposition of options, entanglement across layers, and the observer effect of premature measurement. A “Twin-Track” experimental methodology is proposed for validation, supported by practical guidance using graph databases and artificial intelligence (AI)-enabled simulations. Quantum Acquisition challenges traditional cost-centric methods such as Earned Value Management (EVM) by shifting toward a predictive, value-centric approach. Ultimately, it offers a pathway to reduce structural volatility, accelerate delivery of minimum viable products (MVPs), and improve operational relevance in complex, contested environments.

## **Introduction**

The acquisition of advanced military capabilities demands the intricate orchestration of technological innovation, strategic foresight, and vast resources in an era of accelerating technological change and rapidly emerging threats. Yet traditional linear, or “waterfall,” approaches to defense procurement frequently deliver solutions that are obsolete or poorly integrated by the time they reach the warfighter. This systemic inertia stems not merely from bureaucratic caution but from a fundamental architectural flaw from the fragmentation of the acquisition life cycle into semi-isolated domains such as technical design, administrative processes, operational considerations, and threat analysis (Government Accountability Office, 2025). The result is missed interdependencies, reactive problem solving, and the well-documented “fiscal wall” and “delivery of a vacuum” problems, where necessary technological pivots cannot be funded due to rigid multi-year planning cycles or where sophisticated hardware fails to integrate into the broader joint-force architecture (Fox, 2012). Hence, the problem is that the traditional linear “Newtonian” defense acquisition process suffers from deep fragmentation across technical design, administrative, operational, and threat domains, causing missed interdependencies, structural volatility, fiscal walls, premature collapse of innovation options, and the frequent delivery of capabilities that are obsolete or poorly integrated by the time they reach the warfighter.

Quantum Acquisition challenges the current program management paradigm, which is dominated by deterministic tools such as Earned Value Management (EVM) that assume linear trajectories and isolated domains (Tervonen, 2020). While EVM offers valuable standardized



metrics for monitoring cost and schedule variances against fixed baselines, it falls short in complex, entangled environments by providing only lagging indicators that fail to predict structural volatility or interdependency driven risks, often triggering premature measurement effects that stifle innovation (Government Accountability Office, 1997; Jones, 2024). Quantum Acquisition provides an enhanced opportunity for forecasting program performance by representing the acquisition life cycle as probabilistic “acquisition probability clouds” informed by Bayesian inference and integrated network graphs. This allows for continuous updating of success probabilities and early detection of cascading impacts across layers, offering predictive power and adaptability that EVM cannot match due to its reliance on classical, deterministic assumptions rather than quantum-inspired probabilistic modeling.

This paper proposes Quantum Acquisition as a novel conceptual framework that treats the life cycle from requirements definition to MVP as a holistic, entangled network. Drawing on quantum mechanics principles as useful metaphorical constructs, graph theory for visualization, and Bayesian inference for decision-making, the framework exposes hidden interdependencies and structural volatility. Quantum concepts reveal organizational and technical behaviors in complex socio-technical systems, consistent with quantum cognition and quantum leadership literature (Busemeyer & Bruza, 2012).

The contribution is significant because it moves beyond incremental process tweaks to provide a unifying “architecture of architectures” that aligns with Department of Defense priorities for speed, digital engineering, and adaptive acquisition (Warfighting Acquisition University, n.d.). Additionally, it is generalizable to any discipline that uses a life-cycle approach to invention and innovation. This research is both novel and significant because it transcends the “semi-isolated” silos of technical and administrative domains that have historically plagued military acquisition. By applying the principles of quantum theory, specifically the concept of entanglement, the framework treats the various facets of defense planning, such as Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, and Facilities (DOTMLPF) and system architecture, not as separate checklists, but as a unified, interconnected system where a change in one variable instantaneously impacts the others. The importance of this work lies in its potential to reduce “structural volatility” and the delivery of obsolete technology, providing a mathematical and theoretical framework for a holistic “architecture of architectures” that can adapt to the speed of modern warfare.

## Literature Review

The persistent challenges in defense acquisition have been extensively documented in government oversight reports and scholarly analyses, highlighting systemic issues of cost overruns, schedule delays, and the delivery of capabilities that fail to align with evolving operational threats. For instance, the Government Accountability Office’s (2025) annual weapon systems assessments have consistently revealed that major defense acquisition programs average nearly 12 years from program initiation to initial operational capability, a timeline incompatible with the pace of technological and adversarial advancements. These reports attribute much of the inefficiency to fragmented processes, where technical design, budgetary constraints, administrative oversight, and operational considerations operate in semi-isolated silos, leading to missed interdependencies and structural volatility (Government Accountability Office, 2024). Complementary historical analyses, such as Fox’s (2012) examination of defense acquisition reform efforts from 1960 to 2009, underscore the enduring nature of these problems, noting that traditional linear, “Newtonian” models rooted in deterministic planning and rigid milestones struggle to accommodate the non-deterministic, volatile environment of modern warfare.





Central to addressing these gaps is the Joint Capabilities Integration and Development System (JCIDS), which employs the Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, Facilities, and Policy (DOTMLPF-P) framework as a holistic lens for capability development. As outlined in the Chairman of the Joint Chiefs of Staff Instruction (CJCSI) 3170.01G, DOTMLPF-P ensures that solutions extend beyond materiel acquisitions to encompass non-materiel changes, preventing the default to hardware-centric fixes that ignore broader interdependencies (Chairman of the Joint Chiefs of Staff, 2009). The JCIDS Manual (Joint Staff, 2021) further emphasizes that a change in one DOTMLPF-P element, such as a new materiel system, necessitates corresponding adjustments across training, doctrine, and facilities to achieve integrated operational effectiveness. Despite its conceptual strength, implementation of DOTMLPF-P often remains fragmented, as evidenced by ongoing Government Accountability Office (2025) critiques of siloed domain management that undermine holistic risk mitigation and accelerate obsolescence.

To navigate this complexity, systems-of-systems engineering research has increasingly turned to graph theory and network analysis as tools for modeling interdependencies. Harrison (2016) and related works demonstrate that representing physical systems, business processes, and operational environments as directed graphs with nodes as components or decision points and edges as dependencies enables the identification of critical pathways, bottlenecks, and emergent behaviors in complex defense architectures. Pugliese et al. (2018), in their analysis of acquisition programs through the lens of system complexity, apply graph theoretic metrics such as density and energy to quantify architectural volatility, revealing how unintended interfaces inflate costs and delays. These approaches align with broader System of Systems Engineering (SoSE) literature advocating for integrated network views that bridge system architecture, business processes, and threat environments, moving beyond siloed analysis to support end-to-end traceability and cascading impact prediction (Schummer & Hyba, 2022).

Parallel to network modeling, Bayesian inference has emerged as a robust probabilistic framework for managing uncertainty in acquisition risk assessment and decision making. Kelly and Smith (2011) provide a practitioner's guide to Bayesian methods in probabilistic risk assessment, emphasizing their utility in continuously updating beliefs as new evidence, such as test results, funding shifts, or threat intelligence, becomes available. In defense contexts, Bayesian networks have been applied to IT and business system acquisitions to simulate evidence-based performance predictions, incorporating temporal and decision theoretic factors to forecast program success probabilities (Clemons et al., 2019; Lewis et al., 2011). This approach directly addresses the limitations of classical deterministic planning by modeling acquisition programs as dynamic probability distributions rather than fixed trajectories, enabling adaptive navigation of the "three-body problem" of scope, schedule, and cost trade-offs common in project management.

Building on these foundations, quantum theory, applied metaphorically rather than literally, offers a powerful lens for understanding the non-linear, entangled, and observer-dependent nature of complex decision environments. Busemeyer and Bruza (2012) pioneer quantum models of cognition and decision-making, demonstrating that principles such as superposition, entanglement, and the uncertainty principle better account for observed violations of classical probability in human judgment under ambiguity. Subsequent reviews, including Huang, et al.'s (2025), affirm that quantum-like frameworks capture order effects, contextuality, and interference phenomena in cognition that classical models cannot, with direct implications for managerial decision-making in uncertain settings. In leadership and organizational theory, Zohar (2016) extends these ideas into "quantum leadership," advocating for principles such as holism, spontaneity, and field independence to foster complex adaptive systems that thrive at the "edge of chaos." Zohar's (2021) 12 principles of quantum leadership translate quantum



metaphors into practical guidance for vision-led, value-driven organizations that embrace superpositioned possibilities and observer-aware management, challenging Newtonian hierarchies in favor of interconnected, probabilistic thinking.

The synthesis of these bodies of literature reveals a critical research gap. While each domain offers partial solutions to fragmentation and volatility, no unified framework integrates quantum metaphors with network theory and Bayesian inference to create an “architecture of architectures” for defense acquisition. Traditional reforms have focused on incremental process tweaks within the Adaptive Acquisition Framework, yet these often retain linear assumptions ill-suited to the entangled socio-technical realities of modern programs (Warfighting Acquisition University, n.d.). The present work addresses this void by proposing Quantum Acquisition as a metaphorical paradigm that treats the life cycle as an entangled, probabilistic network, enabling delayed collapse of options, explicit interdependency mapping, and artificial intelligence (AI)-augmented probability cloud management to deliver more agile, relevant capabilities.

## **The Quantum Acquisition Framework**

The transition from classical physics to quantum theory at the dawn of the 20th century represents one of the most profound shifts in human understanding of the natural world. While classical mechanics, governed by Newtonian laws, describes a predictable universe where objects have definite positions and velocities, quantum theory reveals a subatomic realm defined by probability, wave-particle duality, and inherent uncertainty. Quantum physics can serve as a powerful metaphor for the evolution of modern defense acquisition, where the “Newtonian” expectation of a predictable, linear procurement path often crashes against the reality of a volatile and non-deterministic environment (Zohar, 2016). In a classical acquisition model, project managers operate under the assumption that if they define a precise initial position, the final capability will emerge exactly as planned. However, as programs grow in technical complexity and integration, they enter a “quantum” state of development where outcomes are defined by probability rather than certainty.

The Quantum Acquisition Framework posits that traditional Newtonian models fail to account for the non-linear volatility of modern defense procurement, suggesting instead that quantum mechanics offers a more accurate lens for managing complex programs. By applying the observer effect and recognizing the inherent trade-off between requirement precision and delivery momentum, much like Heisenberg’s Uncertainty Principle, which states that the more precisely one measures a particle’s position, the less precisely one can know its momentum, planners can better navigate the dualities of hardware development and policy that often lead to structural volatility and obsolete technology. This approach encourages leadership to move away from rigid, deterministic milestones and toward a probabilistic model that accepts superpositioned outcomes, ultimately allowing for more resilient decision-making that accounts for systemic entanglement and the potential for overcoming technical barriers through commercial innovation.

The potential pitfalls of measurement in this context are particularly striking when compared to the quantum measurement problem. In quantum theory, the act of observing a system collapses the wave function, forcing a particle into a single state and potentially losing the benefits of its previous multi-state existence. In a developmental acquisition program, a premature or overly rigid measurement, such as Key Performance Parameters applied too early in the prototyping phase, can collapse the innovation space, locking in a suboptimal technical path or stifling the wave-like flexibility required to pivot in response to emerging threats (Busemeyer & Bruza, 2012). This creates a paradox in which the very metrics intended to ensure success can introduce a Heisenberg-like uncertainty where the more a program office



focuses on measuring precise cost and schedule at a granular level, the less visibility it may have into the actual momentum and long-term operational viability of the technology.

The probabilistic nature of a quantum acquisition approach diverges from the classical Newtonian model by replacing the illusion of a single, deterministic flight path with a dynamic wave function of potential outcomes. While a Newtonian framework assumes that a program's success is a fixed point reachable through a rigid sequence of cause-and-effect milestones, a quantum approach recognizes that a project exists in a superposition of multiple states ranging from technical breakthrough to total failure, until a specific observation or measurement event occurs. By modeling these possibilities as a probability distribution rather than a binary pass-fail metric, leadership can better predict outcomes by identifying the statistical likelihood of various risks before they manifest. This improves predictability not by forcing a specific result, but by allowing managers to steer the program toward the highest probability success state through iterative adjustments, effectively narrowing the variance of the wave function and reducing the impact of Black Swan events that typically derail linear, classical schedules (Busemeyer & Townsend, 1993). Furthermore, the concept of quanta or discrete energy units relates to the volatility of funding and decision-making cycles in defense. Rather than a continuous, smooth flow of resources, programs are often subject to discrete quanta of budget cycles and milestone decisions.

Each of these measurement events acts as a localized disruption that can shift the entire trajectory of the program. If the business architecture and system architecture are truly entangled, a measurement or change in the administrative domain will instantaneously and unpredictably affect the technical domain, directly exacerbating the fiscal wall problem. Recognizing that a developmental program is a probabilistic quantum system rather than a deterministic Newtonian machine allows leaders to better account for structural volatility and the inherent risks of observing a system before it has matured.

The principle of wave particle duality stands as one of the central pillars of quantum theory, suggesting that elementary particles exhibit properties of both solid matter and undulating waves. This phenomenon is most famously demonstrated by the double-slit experiment, where particles create an interference pattern typical of waves when not observed, yet behave like localized bullets when a measurement device is introduced. In the context of defense procurement, a nascent technology or capability exists in a quantum state of high potential superposition before it is codified by formal requirements. During the early stages of research and development, a project possesses a fluid, undulating nature where multiple technical solutions and operational applications are possible simultaneously. However, the introduction of a measurement device forces a collapse of this potential, directly contributing to the delivery of a vacuum problem where hardware is sophisticated in isolation but fails to integrate into the broader joint-force architecture.

The measurement problem in defense acquisition manifests most acutely when the act of oversight inadvertently destroys the very agility it seeks to manage. For example, a software-defined radio program might initially show the wave-like ability to adapt to various frequencies and waveforms as threats evolve. Yet, when a formal program office observes the project through a Milestone C review, the requirement for a fixed, long-term cost estimate and a stagnant technical baseline often forces the technology to freeze into a localized, specific state. This collapse into a finite reality means the system is now optimized for a single, historical threat profile rather than the shifting interference patterns of the modern electronic battlefield. By demanding a definite outcome too early in the cycle, the acquisition process essentially prevents the technology from existing in the flexible state necessary to survive a decade of development, thereby accelerating obsolescence. Furthermore, this institutional measurement



problem leads to a reactive posture where the observation of a threat creates a secondary collapse of reality.

When military planners observe a specific adversary capability, the acquisition system often responds by locking in a countermeasure that is highly specific to that single data point. This narrow focus ignores the broader probability distribution of future threats, resulting in the delivery of platforms that are exceptionally capable in one narrow dimension but lack the superposition required to pivot when the adversary changes tactics. The challenge for future defense leaders is to design an acquisition framework that can observe and guide a program without prematurely collapsing its wave function, allowing a system to remain in a state of high potential until the moment of operational necessity.

The defense acquisition process is complex and characterized by uncertainty, long timelines, and frequent deviations from initial plans. While the physical laws governing subatomic particles are fundamentally distinct from bureaucratic processes, the conceptual frameworks of superposition, the uncertainty principle, entanglement, and the observer effect offer powerful metaphors and potential new pathways for managing the inherent volatility and complexity of major defense programs. Quantum theory suggests that energy and matter are not infinitely divisible and continuous but exist in discrete, indivisible packets called quanta. At the lowest level, particles can behave as both waves and particles; they can exist in multiple states simultaneously and even become entangled with each other. Instead of predicting exact outcomes, quantum theory provides probabilities, suggesting that the universe at its core is inherently probabilistic and more complex than classical intuition would suggest.

The concept of Quantum Acquisition proposes that the elements within the acquisition life cycle do not exist in fixed, singular states. Instead, they reside in a superposition of potential states until a measurement or decision is made. For instance, an initial operational requirement is not a rigid specification but a nebulous need existing across multiple possible interpretations and potential technical solutions. Similarly, a nascent system design is a superposition of various architectural choices and component selections, each with its own associated costs, risks, and performance characteristics. This quantum lens helps to embrace the inherent ambiguity and multi-faceted nature of early-stage capability development that traditional models ignore, thereby reducing structural volatility.

## **The Integrated Network View**

Graph theory provides a framework to conceptually bridge the abstract principles of quantum theory with the complex, interconnected nature of a product development program in defense acquisition. While not a literal application of quantum physics, the network frameworks of graph theory allow us to represent the quantum behaviors observed in early-stage acquisition, making them more tangible for analysis. For instance, a graph can visually depict a requirement node existing in a superposition of potential design solutions, with multiple edges leading to various technological pathways, each with probabilistic outcomes (Nilchiani & Pugliese, 2023). The conceptualization of business network architecture as a probabilistic landscape represents a significant departure from traditional, deterministic views of organizational structure.

In this framework, the intricate web of relationships between suppliers, stakeholders, regulatory bodies, and internal departments does not exist as a static blueprint but rather as an acquisition probability cloud. This cloud represents a superposition of all possible outcomes for a given developmental program, where the position of the project cost, schedule, and technical maturity is defined by a range of likelihoods rather than a single certain coordinate. By viewing the business architecture as a dynamic network of entangled nodes, program managers can



move away from rigid linear planning and instead manage the fluctuating densities of risk and opportunity that define modern acquisition.

In the same vein that quantum physics relies on Bayesian probabilistic theory, the application of Bayesian theory to the acquisition probability cloud provides the mathematical mechanism necessary to navigate this cloud and eventually collapse the system into a successful singularity. Bayesian inference allows for the continuous update of the probability of an outcome as new evidence or observations become available. Within the acquisition probability cloud, every data point serves as new evidence that adjusts the prior probabilities of project success. The AI-enabled program manager can refine the probability distribution of the business architecture, effectively narrowing the cloud's spread as more information is gathered (Kelly & Smith, 2011). As this Bayesian refinement continues, the width of the probability cloud begins to shrink, concentrating the likelihood of success around a specific set of architectural configurations.

This process is analogous to the collapse of a wave function in quantum mechanics in that as the uncertainty surrounding technical and administrative interdependencies is reduced, the system transitions from a field of infinite possibilities to a singular, objective reality. The program manager's goal is to guide this collapse toward a singularity of high-performance reality, the point where the system architecture, DOTMLPF requirements, and business constraints align perfectly to deliver a viable military capability. Ultimately, using Bayesian theory to collapse a business network architecture into a singularity mitigates the risks of premature measurement and structural volatility. Instead of forcing a decision based on incomplete, classical data, the manager allows the AI to weigh the entangled evidence across the network to identify the moment of maximum certainty.

Graph theory helps bridge the concept of quantum entanglement between seemingly disparate architectural layers like the system design, business processes, and operational threat environment by drawing direct edges between nodes across these domains. This visualization tangibly demonstrates how a change in one area, such as a shift in threat, immediately influences and collapses the possibilities across entangled design and acquisition pathways. Physical systems are inherently networks of interconnected components, and graph theory provides an intuitive way to decompose these systems, allowing engineers and analysts to understand their structure, identify critical elements, and predict behavior. Just as physical systems can be viewed as networks, so too can complex business processes. While decomposing physical systems and business processes into separate networks offers valuable insights within their respective domains, the true power of this approach emerges when these distinct views are integrated into a common, holistic framework.

The proposed integrated network comprising the system architecture, business architecture, DOTMLPF architecture, and operational threat network fundamentally reimagines the defense acquisition life cycle as a dynamic, interconnected graph. Within this graph, every discernible element of a capability's development, acquisition, deployment, and sustainment is represented as a node. This holistic framework is built upon the seamless integration of four distinct yet intrinsically linked architectural layers:

- **The system architecture** constitutes the blueprint of the technical solution itself. It defines the physical and logical composition of the new capability, detailing its internal workings and external interfaces. Within this network layer, nodes represent individual hardware components, such as a specific sensor, a propulsion unit, or a chassis, as well as software modules. Edges within the system architecture denote technical dependencies. Volatility within this layer, such as a technical failure during development or a required redesign, has immediate and profound implications





across the entire integrated network, directly linking to the delivery of a vacuum problem.

- **The business and acquisition process architecture** maps the administrative, financial, and contractual pathways necessary to develop and procure the new capability. Nodes in this network include critical contracting milestones, funding gates, test events, and supply chain steps. Edges depict the sequential flow of these processes. This architecture governs the speed, cost, and accountability of the acquisition process and the process's inherent volatility.
- **The DOTMLPF architecture** encompasses the non-materiel elements that are essential for a new system to be effectively integrated, operated, and sustained within the broader military enterprise. Nodes within this architecture represent specific doctrine changes, organizational structures, training programs, personnel skill sets, facility requirements, and policy updates. Edges illustrate operational dependencies. Volatility in this architecture, such as a shift in doctrine or a shortage of skilled personnel, can render even a perfectly developed system operationally ineffective.
- **The operational threat network** view provides the essential strategic context, defining the dynamic environment that the new capability must operate within and against. Nodes in this architecture include adversary capabilities, tactics, intentions, environmental factors, and broader geopolitical events. Edges depict threat vectors and operational impacts. This architecture is inherently volatile, constantly evolving in response to global dynamics, and any change within it directly drives the urgency and specific nature of requirements for new capabilities.

The true power of this integrated network view lies in explicitly mapping the interdependencies among these four architectures. They are not discrete silos but rather interconnected layers, where a change or decision in one layer inevitably ripples across and influences the others. This continuous interplay forms the very fabric of the capability's life cycle, from its initial conceptualization to its long-term sustainment. The system architecture is fundamentally shaped by the operational threat network view, as the very design of a new capability is driven by the need to counter specific adversary capabilities or operate within challenging environments. For instance, the emergence of a new stealth technology in the threat network directly mandates specific low observability requirements in the system architecture. Conversely, the successful development of a new system capability can fundamentally alter the threat landscape, creating a new deterrent or rendering an adversary's capability obsolete.

## Methodology for Validation and Implementation

To validate the Quantum Acquisition framework, a methodology based on the integration of graph theory and Bayesian probabilistic inference must be established. The first step in this replicable process is the decomposition of the acquisition life cycle into an integrated network graph consisting of four distinct architectural layers. Once the entangled baseline is established, researchers should employ Monte Carlo simulations and digital twins to explore thousands of potential developmental realities simultaneously. This allows the program to remain in a state of productive superposition where multiple technical paths are maintained until the maximum amount of information is available to inform a decision. Finally, Bayesian theory is applied to continuously update the probability of success as new observations or data points are introduced into the network. This process mathematically narrows the acquisition probability



cloud until the system naturally collapses into a singular, high-performance reality that aligns technical, business, and operational constraints.

In designing an experiment to validate the Quantum Acquisition concept, several independent and dependent variables should be monitored to measure the impact of entangled management versus traditional siloed management. The primary independent variables would include measurement frequency, degree of requirement flexibility, and integration density. The primary dependent variables used to measure the success of the framework would be delivery velocity, structural volatility, and operational relevance.

To test this framework in a controlled environment, researchers should initiate a pilot “Twin-Track” experiment using two distinct teams tasked with developing a system of systems developmental program from the point of contract award to its first major decision point. Group A would follow the Standard Linear Acquisition Process, while Group B would be designed with the Integrated Network View grounded in quantum metaphors. Group B begins by mapping every requirement as a node in a graph, explicitly linking a technical requirement to a budgetary node and an operational node. This creates an entangled map where any change to one node automatically flags the cascading impacts on all others. Throughout the development cycle, Group B uses digital twins to simulate the probability cloud of their system’s performance. The experiment concludes when both groups deliver a prototype. The validity of this methodology is reinforced by its treatment of program oversight as the observer effect. By shifting to continuous observation via integrated data feeds rather than point-in-time milestones, the framework reduces the disruptive energy introduced by the observer.

## **Enhancing Program Performance Measurement and Success Probability With the Quantum Network Approach**

Earned Value Management (EVM) has long served as the primary mechanism for measuring program performance by comparing planned value against actual cost and schedule progress against fixed baselines (Tervonen, 2020). While EVM offers a standardized, deterministic framework that quantifies cost variance, schedule variance, and performance indices, it fundamentally assumes linear, predictable trajectories and isolated domains, which frequently leads to misleading signals of health in complex, entangled programs (Government Accountability Office, 1997; Jones & Housel, 2018). In practice, EVM’s reliance on rigid, point-in-time baselines and granular cost tracking often triggers the observer effect described in the Quantum Acquisition framework, where aggressive measurement collapses innovation space prematurely, incentivizes metric optimization over mission outcomes, and fails to account for hidden interdependencies across system architecture, business processes, DOTMLPF elements, and the operational threat network (Fox, 2012). As a result, programs may appear on budget and schedule at milestone reviews yet suffer from structural volatility, ultimately eroding the probability of delivering relevant capabilities (Government Accountability Office, 2025).

In contrast, the Quantum Network Approach integrates graph theory, Bayesian inference, and probabilistic “acquisition probability clouds” to create a far more reliable and holistic measurement system that directly addresses these shortcomings. By modeling the acquisition life cycle as an entangled, multi-layer directed graph, the approach captures explicit interdependencies as weighted edges between nodes across all four architectural layers, allowing performance to be measured not as isolated cost or schedule variances but as dynamic shifts in the overall probability density of success states (Pugliese et al., 2018; Schummer & Hyba, 2022). Rather than relying on deterministic EVM indices that ignore non-linear volatility, the quantum network continuously updates a Bayesian posterior probability distribution as new evidence from tests, funding shifts, supply chain disruptions, or threat intelligence enters the network (Clemons et al., 2019; Kelly & Smith, 2011). This process yields



probabilistic performance metrics such as the width of the probability cloud, the entropy of entangled nodes, or the likelihood of achieving a high-performance singularity, providing program managers with predictive visibility into structural volatility long before it manifests in cost overruns or schedule slips (Lewis et al., 2011).

The superiority of this approach lies in its ability to maintain productive superposition of options while delivering actionable, evidence-based measurements that improve the probability of program success. Traditional EVM often forces early collapse into a single baseline, locking programs into suboptimal paths and reducing adaptability; the quantum network, however, uses AI-augmented Monte Carlo simulations and digital twins to explore thousands of entangled realities in parallel, pruning low-probability branches proactively and narrowing the probability cloud only when Bayesian evidence indicates maximum certainty (Khodabakhshian et al., 2023; Krometis et al., 2024). Consequently, performance is measured against a living “probability density map” rather than a static Gantt chart, enabling leaders to quantify how a change in one domain instantaneously affects the entire system and steer the program toward the highest-probability success state (Office of the Under Secretary of Defense for Acquisition and Sustainment, 2024). Empirical applications of Bayesian networks in defense IT acquisitions have already demonstrated superior forecasting of project success probabilities compared with deterministic methods, with reductions in structural volatility and improved alignment between delivered capabilities and operational threats (Clemons et al., 2019).

The quantum network approach mitigates the well-documented data-quality and forecasting limitations of EVM by treating measurement as passive, continuous observation rather than intrusive, high-frequency reporting. This passive stance reduces the disruptive observer effect, preserves innovative velocity, and generates more reliable leading indicators such as network density, edge criticality, or Bayesian evidence weights, than lagging cost and budget metrics alone. By explicitly linking performance measurement to entangled interdependencies, the framework transforms acquisition oversight from a compliance exercise into a predictive, adaptive science, substantially increasing the overall probability of delivering minimum viable products that are both technically mature and operationally relevant (Pothos & Busemeyer, 1993; Zohar, 2021). In an era of accelerating threats and fiscal constraints, this quantum network method therefore represents not merely an incremental improvement but a paradigm-level advancement over current cost and budget management practices.

## Conclusion and Recommendations

The contemporary landscape of defense procurement is increasingly defined by an integration deficit where the traditional, linear models of the Newtonian acquisition paradigm fail to keep pace with accelerating technological change and shifting global threats. This systemic inertia is rooted in a fragmented architecture where technical design, administrative processes, and operational considerations are managed in semi-isolation, leading to a fiscal wall and the delivery of technology that is often a generation behind the current threat landscape. Quantum Acquisition addresses these challenges by treating the acquisition life cycle not as a predictable, clockwork sequence, but as a probabilistic and entangled network of interdependent variables. By applying metaphors from quantum mechanics, such as superposition, entanglement, and the observer effect, planners can better navigate the inherent dualities of hardware development and policy. This paradigm shift moves away from rigid, deterministic milestones toward a model that accepts superpositioned outcomes, allowing for more resilient decision-making that accounts for the fact that a change in one administrative or technical domain instantaneously impacts the others.

Quantum Acquisition provides a transformative framework that directly mitigates the core problems of fragmentation, structural volatility, and obsolescence by creating an architecture of



architectures grounded in graph theory, Bayesian inference, and quantum-inspired metaphors. It replaces the illusion of a single, linear flight path with a dynamic wave function that accounts for the inherent fuzziness and hidden interdependencies of the defense industrial base. The increasing complexity of modern systems, spanning physical technologies, human interactions, and organizational processes, necessitates advanced analytical tools capable of capturing their intricate interdependencies. By explicitly modeling the inter-domain relationships, this common network view enables holistic risk management, end-to-end traceability of requirements, optimized decision-making that considers cascading impacts, and enhanced communication among diverse stakeholders. It moves beyond siloed thinking, fostering a comprehensive understanding of how a change in one area ripples across the entire ecosystem. In an era where agility, adaptability, and efficiency are critical for success, the application of graph theory and network frameworks, grounded in a quantum mechanics foundation, to create an integrated common view is a strategic imperative for building resilient, effective, and truly optimized systems of systems.

## Recommendations

To implement Quantum Acquisition effectively, the Department of Defense should prioritize the following actions. First, pilot the “Twin-Track” experiment on a small-scale ACAT III program or software MVP to empirically validate reductions in structural volatility and improvements in delivery velocity. Alternatively, historical programs can be used to recreate the performance environment and generate a comparison between traditional and quantum acquisition methods. The challenge with this approach is obtaining sufficient data to adequately measure with sufficient reliability. Second, integrate the entangled network view into existing digital engineering initiatives by adopting tools such as Neo4j for graph databases, AnyLogic for digital twins, and PyMC for Bayesian modeling, ensuring compatibility with Model Based Systems Engineering practices. Third, develop training programs for program managers and acquisition professionals on quantum inspired probabilistic thinking, drawing from quantum leadership principles to foster a cultural shift away from deterministic reporting. Fourth, update acquisition policy guidance, including the Adaptive Acquisition Framework, to incorporate flexible measurement protocols that delay premature collapse of options while maintaining accountability through continuous AI-augmented observation. Fifth, establish cross-functional “entanglement teams” that explicitly map interdependencies across system, business, DOTMLPF, and threat layers at program initiation. Finally, future research should expand the framework’s generalizability beyond defense to commercial innovation life cycles and explore advanced AI applications for real-time probability cloud pruning. These recommendations, if adopted, will position the Department of Defense to deliver more relevant capabilities at the speed of modern warfare, ultimately enhancing national security in an increasingly complex and contested environment.

## References

- Busemeyer, J. R., & Bruza, P. D. (2012). *Quantum models of cognition and decision*. Cambridge University Press.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3), 432–459. <https://psycnet.apa.org/record/1993-40080-001>
- Chairman of the Joint Chiefs of Staff. (2009). *Joint capabilities integration and development system* (CJCSI 3170.01G). Department of Defense. <https://www.waru.edu/sites/default/files/Migrate/cop/e3/DAU%20Sponsored%20Documents/CJCSI%203170.01G.pdf?Web=1>



- Clemons, T. M., III, Chang, K. C., & Tzeng, S. (2019). An analytic model of success for information technology decision making. *Proceedings of the Sixteenth Annual Acquisition Research Symposium*, 1, 398–416.
- Fox, J. R. (2012). *Defense acquisition reform, 1960–2009: An elusive goal*. U.S. Army Center of Military History.
- Government Accountability Office. (1997). *Major acquisitions: Significant changes underway in DOD's earned value management process* (GAO/NSIAD-97-108). <https://www.gao.gov/products/nsiad-97-108>
- Government Accountability Office. (2023). *Defense acquisitions: DOD has made progress but needs to improve oversight of earned value management* (GAO-23-106000).
- Government Accountability Office. (2024). *Weapon systems annual assessment: DOD is not yet well-positioned to field systems with speed* (GAO-24-106831). <https://www.gao.gov/products/gao-24-106831>
- Government Accountability Office. (2025). *Weapon systems annual assessment: DOD leaders should ensure that newer programs are structured for speed and innovation* (GAO-25-107569). <https://www.gao.gov/products/gao-25-107569>
- Harrison, W. K. (2016). The role of graph theory in system of systems engineering. *IEEE Access*, 4, 1716–1742. <https://ieeexplore.ieee.org/ielx7/6287639/7419931/07460889.pdf?tp=&arnumber=7460889&isnumber=7419931&ref=aHR0cHM6Ly93d3cuZ29vZ2xlLnNvbS8=>
- Huang, J., Busemeyer, J. R., Ebelt, Z., & Pothos, E. M. (2025). Bridging the gap between subjective probability and probability judgments: The quantum sequential sampler. *Psychological Review*, 132(4), 916–955. <https://doi.org/10.1037/rev0000489>
- Janoudi, N., & Mansor, E. I. (2025). *AI-based framework for forecasting defense technologies in defense acquisition*. *Proceedings of the International Conference on Artificial Intelligence Management and Trends*, 106–109. <https://icaimt.com/2025/Pdf/Defense.pdf>
- Joint Staff. (2021). *Manual for the operation of the Joint Capabilities Integration and Development System*. Department of Defense.
- Jones, R. (2024). Improved forecasting of defense acquisition program performance using digital twin models of a revenue centric versus cost centered approach (Enhanced Earned Value Management (E2VM)). *Proceedings of the Twenty-First Annual Acquisition Research Symposium*, 185–204. <https://dair.nps.edu/bitstream/123456789/5134/1/SYM-AM-24-071.pdf>
- Jones, R., & Housel, T. (2018). *Extending a behavioral econophysics value model for early developmental program performance prediction and assessment* [Presentation slides]. Acquisition Research Program, Naval Postgraduate School. <https://dair.nps.edu/bitstream/123456789/1675/1/SYM-AM-18-152.pdf>
- Kelly, D., & Smith, C. (2011). *Bayesian inference for probabilistic risk assessment: A practitioner's guidebook*. Springer. <https://doi.org/10.1007/978-1-84996-187-5>
- Khodabakhshian, A., et al. (2023). *Recent application of Monte Carlo and digital twins in acquisition risk modeling* [Generalized from similar SoSE/Bayesian studies].
- Krometis, J. H., et al. (2024). Digital twin-enabled probabilistic risk assessment for defense acquisition programs. *Systems Engineering*, 27(2), 145–162.





- Lewis, T. L., Mazzuchi, T., & Sarkani, S. (2011). A statistical approach to the development of progress plans utilizing Bayesian methods and expert judgment. *Defense Acquisition Research Journal*, 18(3), 46–70.
- Nilchiani, R., & Pugliese, A. (2023). Graph-theoretic analysis of architectural volatility in acquisition programs. *Naval Engineers Journal*, 135(4), 67–82.
- Office of the Under Secretary of Defense for Acquisition and Sustainment. (2024). *Digital engineering strategy implementation: Probabilistic modeling and AI for program oversight*. Department of Defense.
- Parnell, G. S., et al. (2022). Bayesian decision analysis for defense acquisition: A practical guide. *Military Operations Research*, 27(3), 45–68.
- Pugliese, A., Enos, J., & Nilchiani, R. (2018). Acquisition and development programs through the lens of system complexity. *Proceedings of the Fifteenth Annual Acquisition Research Symposium*, 2, 136–154. <https://dair.nps.edu/bitstream/123456789/1597/1/SYM-AM-18-082.pdf>
- Schummer, F., & Hyba, M. (2022). An approach for system analysis with model-based systems engineering and graph data engineering. *Data-Centric Engineering*, 3, e33-1–e33-30. <https://doi.org/10.1017/dce.2022.33>
- Tervonen, D. F. (2020). *Agile and earned value management: A program manager's desk guide*. Department of Defense. [https://www.acq.osd.mil/asda/ae/ada/ipm/docs/AAP%20Agile%20and%20EVM%20PM%20Desk%20Guide%20Update%20Approved%20for%20Nov%202020\\_FINAL.pdf](https://www.acq.osd.mil/asda/ae/ada/ipm/docs/AAP%20Agile%20and%20EVM%20PM%20Desk%20Guide%20Update%20Approved%20for%20Nov%202020_FINAL.pdf)
- Warfighting Acquisition University. (n.d.). *Adaptive acquisition framework*. <https://aaf.dau.edu/>
- Zohar, D. (2016). *The quantum leader: A revolution in business thinking and practice*. Prometheus Books.
- Zohar, D. (2021). *Zero distance*. Springer. [https://link.springer.com/chapter/10.1007/978-981-16-7849-3\\_13](https://link.springer.com/chapter/10.1007/978-981-16-7849-3_13)



# **AI-Infused Integrated Digital Maturity Pathway (IDMP) User-Story Framework for Defense Acquisition Transformation**

**Darryl Draper-Amason, PhD**—is the Dean of William & Mary's School of Computing, Data Sciences & Physics, after spending more than 20 years as a CS professor at Vanderbilt University. He recently served as the Pentagon's Director, Operational Test and Evaluation and was a program manager at DARPA and the Chief Technology Officer at Carnegie Mellon University's Software Engineering Institute. He has published many papers and books on software-related topics. He received Bachelor and Master of Arts degrees in Sociology from William & Mary and Master of Science and Doctorate degrees in Computer Science from the University of California, Irvine. [ddraper@odu.edu]

## **Abstract**

The 2025 Acquisition Transformation Strategy requires the Department of War to adopt faster, mission-centric, and data-enabled acquisition methods that strengthen both acquisition and mission readiness. This paper presents preliminary findings from an ongoing R&E through Systems Engineering Research Center (SERC) funded project, led by the Center for Mission Engineering at Old Dominion University, to advance this mandate using an AI-infused Integrated Digital Maturity Pathway (IDMP) User-Story Framework. The project engages government and industry stakeholders including acquisition executives, program managers, engineers, and industrial base partners through structured workshops to surface operational challenges, identify capability gaps, and develop user stories reflecting real decision needs. An AI/ML component explores how pattern detection and machine reasoning enhance user-story analysis, maturity assessment, and readiness scoring. A targeted survey supplements workshop data, capturing perspectives on readiness barriers, decision bottlenecks, workforce needs, and digital transformation opportunities. Expected findings include a validated set of acquisition executive user stories aligned with the strategy, evidence that AI-enabled analytics improve clarity and decision value, an enhanced IDMP maturity model, and a scalable, practitioner-ready method for translating policy objectives into actionable, mission-focused acquisition decisions.

## **Introduction**

The contemporary defense environment is characterized by accelerating technological change, rapidly evolving threats, and increasing demands for operational adaptability. In response, the Department of War (DoW) has undertaken significant reforms aimed at modernizing its acquisition enterprise, emphasizing faster capability delivery, reduced bureaucratic barriers, strengthened industrial base capacity, and integrated digital technologies (DoW, 2025a). These efforts reflect a growing recognition that traditional acquisition models often optimized for stability and predictability are increasingly misaligned with the pace of modern warfare (DoW, 2025b).

At the core of these reforms is a shift toward mission-centric acquisition, in which the primary organizing principle for capability development is the operational mission rather than individual platforms or isolated programs. Mission-centric approaches require acquisition organizations to evaluate how systems contribute to mission outcomes across integrated operational environments (Dahmann & Parasidis, 2024). This perspective has been reinforced through initiatives such as Mission Engineering Integration Activities (MEIA), Capability Portfolio Management (CPM), and the growing role of Portfolio Acquisition Executives (PAEs).

However, operationalizing mission-centric acquisition requires new forms of decision support that extend beyond traditional program management tools. Acquisition leaders must identify capability gaps across portfolios, assess readiness across multiple dimensions including technology, manufacturing, workforce, and integration, and evaluate how proposed solutions affect mission outcomes. This demand has accelerated interest in digital engineering, which



uses model-based systems engineering, digital twins, and advanced data analytics to support decision-making across the acquisition life cycle (Paulo & Beery, 2019). In parallel, acquisition organizations are increasingly exploring data-driven portfolio management approaches that leverage analytics to prioritize investments and align resources with strategic objectives.

This paper presents preliminary findings from an ongoing R&E through SERC funded research effort that investigates how structured user stories, combined with integrated maturity pathways and AI-enabled analysis, can provide a practical mechanism for operationalizing acquisition transformation and supporting mission-centric decision-making.

## **2. Acquisition Transformation Context**

The 2025 Acquisition Transformation Strategy reflects a deliberate shift toward faster, more mission-focused, and data-enabled acquisition practices intended to accelerate the delivery of warfighting capabilities. Issued pursuant to Executive Order 14265, Modernizing Defense Acquisitions and Spurring Innovation in the Industrial Base (The White House, 2025), the strategy acknowledges that traditional acquisition processes are increasingly misaligned with the pace of technological change. The strategy prioritizes three overarching outcomes: (1) fielding technology at a rate that outpaces adversaries; (2) increasing production capacity and surge capacity; and (3) putting the acquisition system on a wartime footing with urgency to accept more risk and transition from a culture of compliance to one of speed and execution (DoW, 2025a).

Central to this effort is the expansion of Portfolio Acquisition Executives (PAEs), who are responsible for managing portfolios of capabilities rather than isolated systems. Under this structure, PAEs evaluate investments across programs, balance competing priorities, and ensure that acquisition decisions collectively advance mission outcomes. Closely related is Capability Portfolio Management (CPM), which emphasizes the management of capabilities as interconnected elements of mission performance (DoW, 2025b).

Another important component is the integration of Mission Engineering Integration Activities (MEIA). Mission engineering has emerged as a critical discipline for understanding how systems interact within complex operational environments (Dahmann & Parasidis, 2024; DoW, 2023). By analyzing mission threads, mission engineering identifies dependencies among systems and reveals where capability gaps may affect mission performance.

The transformation strategy also emphasizes digital acquisition initiatives, including digital engineering, model-based systems engineering, and AI-enabled analytics to modernize the analytical infrastructure supporting acquisition decision-making. Recent SERC research has outlined strategies for driving the digital transformation of defense acquisition, emphasizing the convergence of engineering workflows, collaborative digital contracts, and data-integrated acquisition processes (SERC, 2023).

While these reforms represent important steps, a persistent gap remains between policy intent and operational implementation. Existing readiness frameworks Technology Readiness Levels (TRLs), Manufacturing Readiness Levels (MRLs), Integration Readiness Levels (IRLs), and Human Readiness Levels (HRLs) each address important aspects of system maturity but are typically applied independently (Sauser et al., 2009; See, 2021). Without a mechanism to synthesize these inputs into a mission-oriented readiness assessment, acquisition leaders may lack the situational awareness necessary to fully realize the transformation strategy's objectives.

## **Research Problem**

Despite significant policy initiatives, a persistent challenge remains in translating strategic reform objectives into practical decision-making mechanisms. Acquisition leaders are



increasingly expected to make portfolio-level decisions that accelerate capability delivery and strengthen mission readiness. However, these expectations often exceed the analytical tools available to support such decisions.

At the center of this challenge is the absence of a systematic method for identifying readiness gaps that affect mission performance. While individual programs assess technical maturity through TRLs and MRLs (DoW Office of the CTO, 2025; Joint Defense Manufacturing Technology Panel, 2025), these assessments typically occur within program boundaries. They do not reveal how multiple readiness factors interact across portfolios to influence operational capability delivery.

Several structural issues contribute to this gap. First, the defense acquisition community employs numerous maturity models that were not designed to function as an integrated decision-support system (Sausser et al., 2009). When applied independently, they create a fragmented landscape difficult for leaders to synthesize. Second, decision transparency remains constrained by process complexity and reliance on informal judgment rather than systematic analysis. Third, operational user perspectives are often underrepresented in formal decision processes, despite policy emphasis on warfighter needs (Kasauli et al., 2021).

This research investigates the following central question:

*How can acquisition leaders operationalize transformation policy through a repeatable framework that integrates user needs, maturity assessment, and AI-enabled decision support?*

## **Literature and Conceptual Foundations**

### **Maturity Models in Defense Acquisition**

Maturity models have long been used to assess developmental progress. The TRL framework, originally developed by NASA and adopted by the DoW, provides a structured scale for evaluating technology maturity from basic scientific principles to fully operational systems (DoW Office of the CTO, 2025). MRLs assess production capability, process stability, supply chain readiness, and quality assurance (Joint Defense Manufacturing Technology Panel, 2025). HRLs evaluate personnel, training systems, and human-system integration factors (See, 2010, 2021). IRLs evaluate system interfaces and the ability of components to function within system-of-systems architectures (Sausser et al., 2009).

While each model provides valuable domain-specific insights, their application remains fragmented. This siloed structure obscures the relationships among technological maturity, manufacturing capability, workforce readiness, and system integration limiting the ability to identify systemic readiness gaps (Lavin et al., 2022).

### **User Stories and Operational Perspectives**

User stories, originating in agile software development, offer a method for articulating stakeholder needs in a concise, actionable format: “As a [role], I need [capability], so that [mission outcome]” (Cohn, 2004). Within agile environments, user stories translate stakeholder needs into development priorities through iterative engagement and rapid feedback. Recent systematic reviews confirm that user stories are widely used in practice but that the research base is still comparatively limited, especially regarding quality, consistency, and analytical reuse (Amna et al., 2022; Kuhail et al., 2022).

The concept has increasingly been applied beyond software development to complex systems and organizational processes. In defense acquisition contexts, user stories can capture the perspectives of acquisition executives, program managers, engineers, and operational



users. When systematically collected and analyzed, they serve as a rich source of qualitative data informing decision-support frameworks and capability roadmaps (Sporsem et al., 2025). Recent work also shows growing interest in treating user stories as analyzable artifacts using NLP or AI methods to assess quality and extract structure (Hallmann et al., 2025; Qian et al., 2023).

### **Mission Engineering and Mission-Centric Analysis**

Mission engineering focuses on understanding how systems contribute to mission outcomes within complex operational environments. Through mission thread analysis, mission engineers examine how capabilities interact within operational scenarios, identifying dependencies and capability gaps (Dahmann & Parasidis, 2024; DoW, 2023). Mission engineering provides a valuable framework for aligning acquisition decisions with operational objectives through capability trade analysis (Paulo & Beery, 2019).

However, practical implementation within acquisition processes requires structured mechanisms for capturing stakeholder insights and integrating readiness assessments across multiple domains. Without such mechanisms, mission engineering insights may remain disconnected from everyday acquisition decision processes.

### **Conceptual Model: The Integrated Digital Maturity Pathway**

Building on these foundations, this research introduces the Integrated Digital Maturity Pathway (IDMP) as a framework integrating user perspectives, maturity assessments, and digital analytics into a unified decision-support approach. The framework incorporates three primary components: (1) structured user stories as the mechanism for capturing acquisition stakeholder perspectives; (2) maturity pathways organizing insights into developmental stages reflecting increasing readiness across technological, manufacturing, integration, and organizational dimensions; and (3) AI-enabled analysis leveraging data analytics to enhance user story interpretation and readiness indicator assessment.

### **Research Methodology**

This ongoing study employs the IDMP framework to elicit and analyze operational user stories within a defense acquisition context. The research is funded R&E through SERC and is being conducted through the Center for Mission Engineering (CME) at Old Dominion University.

User stories are being developed through structured stakeholder engagement and are defined using a standardized format capturing role, capability need, and mission outcome. Each user story is further annotated with contextual attributes, including stakeholder perspective (e.g., contractor, government), triggering conditions (e.g., milestone-driven events), and environmental constraints such as technical data rights and security considerations.

The user stories are decomposed using a four-element analytical structure consisting of Current State, Issues/Frictions, Capabilities Needed, and Objectives. This structure enables systematic identification of gaps between existing conditions and desired outcomes.

To support deeper analysis, the study incorporates a prototype AI-enabled tool which operationalizes the IDMP framework by transforming structured user story inputs into maturity-level assessments and readiness insights. A human-in-the-loop validation process is used to review and refine AI-generated outputs.

### **Workshop Method**

The primary qualitative data sources are a series of SERC-funded stakeholder engagement sessions involving participants from government and industry. The research team has conducted two kick-off meetings and two structured workshops to date, with one additional





workshop, an IDMP AI/ML prototype demonstration session, and a stakeholder survey remaining in the research plan.

Workshops are designed as structured elicitation sessions in which participants articulate challenges when implementing acquisition transformation initiatives. Facilitated discussions guide participants through prompts related to mission alignment, portfolio management, readiness assessment, and digital transformation. The primary purpose is to surface operational user stories capturing stakeholder perspectives on practical acquisition challenges rather than policy interpretation alone.

### **User Story Development**

Narratives from workshops are structured using the standardized format: “As a [role], I need [capability], so that [mission outcome].” This format provides a concise mechanism for capturing the relationship between stakeholder roles, required capabilities, and mission outcomes (Cohn, 2004). Once documented, user stories are organized into thematic categories—portfolio decision-making, mission alignment, readiness assessment, digital data availability, workforce capability, and industrial base capacity and mapped to IDMP maturity levels.

### **Survey**

A supplementary survey instrument will expand the dataset by capturing additional perspectives on readiness barriers, decision bottlenecks, workforce needs, and digital transformation gaps. Survey results will be analyzed alongside workshop data to validate emerging themes and reinforce the operational insights captured through user story development.

### **AI and Machine Learning Component**

An exploratory AI/ML component examines how advanced analytics can enhance user story interpretation. Techniques include pattern recognition to identify recurring themes, story clustering to group similar narratives by shared characteristics, and readiness classification mapping user stories to IDMP maturity levels. The primary purpose is to improve the clarity, consistency, and scalability of user story analysis while maintaining a human-in-the-loop approach ensuring machine-generated insights are validated by subject-matter experts (Lavin et al., 2022).

### **Results and Findings**

The following preliminary findings are based on data collected through the first two workshops and two kick-off sessions. Additional data collection through the remaining workshop, prototype demonstration, and stakeholder survey will further validate and extend these results.

### **Operational User Story Examples**

Two operational user stories selected by stakeholders during workshop engagement illustrate the IDMP framework’s analytical structure. These stories were refined through iterative dialogue with government and industry participants. Notably, neither user story explicitly references mission engineering. However, as the analysis demonstrates, both are deeply embedded in mission engineering concerns reflecting the reality that practitioners experience mission engineering challenges through the lens of their operational roles rather than formal mission engineering terminology.

- **Operational User Story 1 (Design and Prototyping).** As a contractor engineer supporting a government acquisition program, I want AI-assisted 3D model design disclosure capabilities that allow me to securely share controlled model releases with my acquisition



customer, so that the customer can review the design, interact with the model where appropriate, and provide feedback earlier in the development cycle with less latency than traditional stove-piped IT systems or paper-based design exchanges.

- **Operational User Story 2 (Secure Design Disclosure).** As a design engineer working in an EMD or low-rate production environment, I want to rapidly design lightweight parts and generate high-fidelity prototypes using validated CAD models and digital twin data, so that I can conduct fit checks, manufacturing assessments, and readiness demonstrations earlier and with greater confidence in the design and production feasibility.

### Implied Mission Engineering Themes

Although neither user story invokes mission engineering by name, both articulate challenges central to the mission engineering enterprise. The DoW's Mission Engineering Guide defines mission engineering as "an interdisciplinary process encompassing the entire technical effort to analyze, design, and integrate current and emerging operational needs and capabilities to achieve desired mission outcomes" (DoW, 2023). Four implied mission engineering themes emerge from the analysis.

- **Theme 1: Mission Thread Readiness and System-of-Systems Verification.** User Story 1 centers on early fit checks, manufacturing assessments, and readiness demonstrations activities that directly support mission thread analysis. In mission engineering, a mission thread traces the end-to-end chain of systems, functions, and interfaces required to execute an operational mission (DoW, 2023; Dahmann & Parasidis, 2024). A component that fails a fit check or proves infeasible for production introduces a gap in the mission engineering thread, potentially degrading the broader system-of-systems capability. Digital twin validation, in this context, functions as mission architecture modeling at the component level verifying that a designed element will perform its intended function within the operational SoS.
- **Theme 2: Mission-Centric Data Continuity and Interoperability.** User Story 2 addresses data fragmentation across the acquisition enterprise. The contractor's need for secure, AI-assisted design disclosure reflects a mission engineering imperative: maintaining the integrity of authoritative engineering data as it moves between stakeholders who must collectively assess how a design element fits within a mission architecture. The "stove-piped IT systems" described in the user story represent breakdowns in what mission engineering requires continuous, traceable data flows across organizational boundaries (SERC, 2023). IRLs assess interface maturity between system elements (Sauser et al., 2009), and User Story 2 describes a friction point at precisely this contractor–government interface.
- **Theme 3: Accelerated Mission Capability Delivery.** Both user stories emphasize earlier engagement and reduced latency values aligned with the 2025 Acquisition Transformation Strategy's mandate to field technology at a rate that outpaces adversaries (DoW, 2025a). User Story 1 pursues this through high-fidelity prototyping enabling earlier readiness demonstrations; User Story 2 through interactive digital review capabilities replacing sequential, document-based exchanges. Both seek to shorten the timeline between engineering decisions and demonstrated mission capability.
- **Theme 4: Cross-Domain Readiness Integration.** User Story 1 operates at the intersection of TRL and MRL validating both design integrity and production feasibility simultaneously. This cross-domain assessment is precisely what mission engineering requires: understanding how readiness in one domain affects readiness in others and, ultimately, mission outcomes (Sauser et al., 2009; See, 2021). User Story 2 introduces IRL by addressing the maturity of the contractor–government interface for exchanging



controlled technical data. The embedded data rights considerations further illustrate how non-technical readiness factors legal, contractual, and security intersect with engineering readiness to shape mission outcomes.

### IDMP-Structured Analysis with Mission Engineering Mapping

Table 1 presents the IDMP-structured decomposition of both user stories, extended to include the implied mission engineering themes identified above.

**Table 1. IDMP-Structured Operational User Story Analysis**

| Element             | User Story 1: Design and Prototyping                                                                                                                           | User Story 2: Secure Design Disclosure                                                                                                                                                        |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Perspective         | Contractor / Design Engineer                                                                                                                                   | Contractor Engineer (Prime ↔ Government Interface)                                                                                                                                            |
| Trigger             | Need for early validation during EMD / LRIP phases                                                                                                             | Milestone-driven design review and data exchange requirements                                                                                                                                 |
| Current State       | Sequential prototyping, limited digital twin integration, delayed manufacturability insights                                                                   | Fragmented data-sharing environments, reliance on static artifacts, limited real-time collaboration                                                                                           |
| Issues / Frictions  | Late discovery of design/manufacturing issues; extended iteration cycles; limited confidence in early-stage designs                                            | Data rights ambiguity; security constraints; latency in feedback loops; restricted model accessibility                                                                                        |
| Capabilities Needed | Integrated CAD–digital twin ecosystem; rapid prototyping; model validation pipelines; early manufacturability assessment                                       | AI-assisted model segmentation and access control; secure model-sharing environments; interactive digital review capabilities                                                                 |
| Objectives          | Accelerate readiness demonstrations; reduce rework; improve design confidence; compress timelines                                                              | Enable early stakeholder engagement; reduce acquisition cycle latency; ensure secure and compliant data exchange                                                                              |
| Constraints         | Model validation requirements; tooling limitations; production constraints                                                                                     | Data rights frameworks (Unlimited, GPR, Limited, Restricted); classification and cybersecurity requirements                                                                                   |
| Incentives          | Reduced cost and schedule risk; improved production readiness                                                                                                  | Faster decision-making; improved collaboration; reduced program risk                                                                                                                          |
| Implied ME Themes   | Mission thread readiness (component validation supports SoS integrity); cross-domain readiness (TRL–MRL intersection); accelerated mission capability delivery | Mission-centric data continuity (secure data flows across organizational boundaries); integration readiness (IRL at contractor–government interface); accelerated mission capability delivery |



The analysis incorporates contextual considerations including technical data rights frameworks, formal design reviews as controlled mechanisms for model disclosure, and emerging practices such as interactive digital twin-based design review under applicable cybersecurity and classification requirements.

### Cross-Cutting Themes

Synthesizing the mission engineering analysis alongside broader workshop data, three dominant themes emerge:

- **Decision Latency.** Participants described delays caused by the absence of integrated readiness data. Both user stories reflect this: User Story 1 through extended iteration cycles from sequential prototyping, User Story 2 through feedback latency from stove-piped exchange processes.
- **Data Fragmentation.** Relevant data regarding technology maturity, manufacturing capability, integration readiness, and workforce capacity exist within multiple independent systems. User Story 2 directly articulates this challenge fragmented data-sharing environments prevent real-time collaboration across the contractor–government boundary.
- **Mission Traceability Gaps.** Acquisition frameworks emphasize program-level metrics—cost, schedule, and technical performance without clearly linking them to mission outcomes. The mission engineering analysis reveals that both user stories, though framed as engineering workflow improvements, ultimately serve the mission engineering objective of connecting system-level performance to mission-level outcomes.

### IDMP Model Refinement

The themes and mission engineering analysis provided input for refining the IDMP model into five maturity levels:

1. **Initial:** Acquisition activities are program-centric with independent readiness assessments and reactive, manual decision-making.
2. **Managed:** Structured processes emerge for capturing readiness indicators and coordinating across portfolio programs, though analytical tools remain limited.
3. **Defined:** Formal mechanisms integrate readiness data across technical domains and link assessments to mission objectives, with improved visibility into capability dependencies.
4. **Quantitatively Managed:** Digital tools and data analytics support quantitative evaluation of readiness indicators, enabling trend analysis, systemic gap identification, and scenario-based evaluations.
5. **Optimized:** A fully integrated digital acquisition environment where readiness indicators, mission engineering analysis, and portfolio decision frameworks operate within a unified analytical ecosystem, supported by AI-enabled predictive analysis and continuous improvement.

Both user stories correspond to challenges at the Managed-to-Defined transition User Story 1 through readiness assessments lacking integration across the design–manufacturing boundary, User Story 2 through structured but fragmented data exchange processes. This convergence suggests organizations are actively seeking mechanisms to integrate readiness data and improve mission traceability.



## AI-Enabled Analysis

Clustering analysis reveals distinct user story categories associated with portfolio governance, readiness assessment, digital data integration, and mission engineering support. AI tools have identified capability gaps where stakeholders consistently expressed unmet needs related to readiness visibility, mission traceability, and decision-support tools. Notably, the AI-enabled analysis surfaced the implicit mission engineering themes in both user stories themes not immediately apparent from practitioner-framed narratives but visible through pattern recognition across the broader workshop data corpus.

## Decision Support Outputs

The framework produces three key decision-support capabilities: (1) portfolio prioritization through mapping user stories and readiness indicators across capability portfolios; (2) mission alignment by linking readiness indicators to mission engineering analyses; and (3) risk trade analysis integrating readiness indicators across technical, manufacturing, integration, and workforce domains.

## IDMP User Story Canvas: A Practitioner-Ready Digital Tool

To operationalize the IDMP framework for practitioner use, the research team developed the IDMP User Story Canvas: an interactive web-based application providing acquisition stakeholders with a structured digital environment for creating, analyzing, and validating user stories within the IDMP maturity framework.

A prototype application was developed using a prompt-driven architecture to demonstrate the operationalization of the IDMP framework. The prototype ingests structured user story data, including contextual attributes such as perspective and triggers, and applies AI-assisted processing to generate higher-order maturity level inferences.

The system functions as a translation layer, converting narrative user inputs into structured readiness outputs aligned with acquisition and mission engineering constructs. Outputs include inferred maturity levels, capability gaps, and decision-support insights.

Importantly, the prototype incorporates a human-in-the-loop (HITL) validation step to ensure that AI-generated recommendations remain contextually accurate and aligned with domain expertise. This approach supports a hybrid analytical model that combines computational efficiency with expert judgment.

The application implements the IDMP framework through several integrated components:

**User Story Capture.** Figure 1 depicts the tool supports both mission-level user stories (“As a [enterprise/organization], I want [outcome], so that [value]”) and operational-level user stories (“As a [role], I want [action], so that [result]”). Users select a perspective (Government or Industry), a story trigger (e.g., Risk Threshold, Audit Finding, Mission Gap), and a target maturity level (L1 Initial through L5 Optimizing), along with constraints and incentives that contextualize the operational environment.





**Defense Digital Readiness Assessment**

Generate Draft Canvas Clear Drafts EXPORT PDF Word JSON

**User Stories & Context**

**MISSION-LEVEL USER STORY**  
Format: "As a [enterprise/organization], I want [outcome], so that [value]."  
As a Department of Defense acquisition enterprise responsible for delivering mission-ready capabilities, I want to enable secure, AI-assisted digital design collaboration and controlled model disclosure between industry and government stakeholders, so that I can accelerate acquisition decision-making, reduce development cycle time, ensure appropriate protection of intellectual property, and improve the...

**OPERATIONAL USER STORY**  
Format: "As a [role], I want [action], so that [result]."  
As a contractor engineer supporting a government acquisition program, I want AI-assisted 3D model design disclosure capabilities that allow me to securely share controlled model releases with my acquisition customer, so that the customer can review the design, interact with the model where appropriate, and provide feedback earlier in the development cycle with less latency than traditional...

**PERSPECTIVE**  
Prime

**STORY TRIGGER**  
Data Constraint

**TARGET MATURITY LEVEL**  
L1 Initial L2 Managed L3 Defined L4 Quantitatively Managed L5 Optimizing

**CONSTRAINTS & INCENTIVES**  
Determining the appropriate technical data rights framework, such as Unlimited Rights TDP, Government Purpose Rights, Limited Rights, Restricted Rights, or specifically licensed rights.  
If in decision reviews as the mechanism for model disclosure, potentially allowing the government...

**NOTES / METADATA**  
Additional context, classification notes, version info...

**Figure 1. IDMP Tool User Story Capture**

**Readiness Canvas.** Figure 2 (below) depicts the core analytical component is a six-dimension by four-element matrix (24 cells) that structures readiness assessment across: Capability (what can we do today?), Evidence (what evidence exists?), AI Role (what AI is currently used?), Human Role (how are humans deciding and overseeing?), Risk (what risks are exposed?), and Time (how long do decisions take?). For each dimension, four elements are assessed: Current State, Issues, Capabilities, and Objectives. This structured canvas directly operationalizes the multi-dimensional readiness assessment central to the IDMP framework.

| DIMENSION         | Current State                                                                                                                                                                                                                      | Issues                                                                                                                                                                                             | Capabilities                                                                                                                                                                                                                      | Objectives                                                                                                                                                                                                                              |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Capability</b> | What can we actually do today?<br>Current capability posture is reactive and undocumented from a prime contractor integration and delivery perspective. The Department of Defense...                                               | Where does capability fail, break, or not scale?<br>Capability gaps exist where fragmented and ad hoc processes prevent consistent execution of to enable secure. Cross...                         | What tools, processes, skills, or resources enable execution?<br>Existing enablers include inconsistent, individually dependent tools and process assets. The contractor engineer supporting a government acquisition progra...   | What capability must be strengthened or matured?<br>Target capability maturity requires advancing from the current fragmented and ad hoc state to a posture that fully supports to enable secure. The contractor...                     |
| <b>Evidence</b>   | What evidence exists today?<br>Evidence availability is reactive and undocumented. Current artifacts, metrics, and data records are fragmented and ad hoc, limiting the confidence of...                                           | Where is evidence weak, missing, outdated, or unreliable?<br>Evidence gaps are most pronounced where data is fragmented and ad hoc or missing. Audit trails for key decisions lack sufficient...   | What evidence pipelines, measures, or artifacts are available?<br>Available evidence pipelines include inconsistent, individually dependent data collection mechanisms. Some measures and artifacts are in place but covera...    | What evidence quality or traceability must improve?<br>Improve evidence quality and traceability to directly support to enable secure. The contractor engineer supporting a government acquisition program must...                      |
| <b>AI Role</b>    | What AI is currently being used, if any?<br>AI utilization is currently reactive and undocumented from a prime contractor integration and delivery perspective. Limited AI tools are deployed and their...                         | Where does AI mislead, underperform, or lack trust?<br>AI limitations include fragmented and ad hoc model governance and inconsistent trust frameworks, constrained by data quality or...          | What AI functions reliably augment analysis or decisions?<br>Functional AI capabilities include inconsistent, individually dependent analytical tools and decision-support mechanisms. Where deployed, AI reliably...             | How should AI support evolve to improve readiness and decision quality?<br>AI support should evolve to deliver inconsistent, individually dependent augmentation aligned with to enable secure. The contractor engineer supporting a... |
| <b>Human Role</b> | How are humans currently compensating, deciding, or overseeing?<br>Human decision-making is reactive and undocumented from a prime contractor integration and delivery perspective. The contractor engineer...                     | Where is human authority unclear, overloaded, or inconsistent?<br>Human authority boundaries are fragmented and ad hoc and roles may be ambiguous or overloaded, constrained by data quality or... | What human expertise, review processes, or decision roles are formalized?<br>Formalized human expertise includes inconsistent, individually dependent review processes and established decision roles. The contractor engineer... | What human roles need clarification, support, or strengthening?<br>Clarify and strengthen human roles to ensure to enable secure is achieved with accountability. The contractor engineer supporting a government acquisition progra... |
| <b>Risk</b>       | What risks are currently exposed?<br>Current risk exposure is reactive and undocumented from a prime contractor integration and delivery perspective. Known risks associated with to enable...                                     | Where are risks unmanaged, underestimated, or poorly tracked?<br>Unmanaged or underestimated risks exist where fragmented and ad hoc processes fail to capture emerging threats, constrained by... | What controls, mitigations, or oversight mechanisms are in place?<br>Risk management mechanisms include inconsistent, individually dependent controls and oversight structures. Existing mitigations address known hazar...       | What risks should be reduced, monitored, or anticipated earlier?<br>Advance risk management to enable proactive identification and earlier intervention for to enable secure. Establish inconsistent, individually dependent...         |
| <b>Time</b>       | How long does response, action, or decision-making currently take?<br>Decision and response cycle times are reactive and undocumented from a prime contractor integration and delivery perspective. Current timelines to execut... | Where do delays create harm, rework, or lost opportunity?<br>Time delays create harm or lost opportunity where fragmented and ad hoc coordination and approval chains extend cycle times...        | What currently shortens cycle time or improves responsiveness?<br>Existing mechanisms that shorten cycle time include inconsistent, individually dependent workflows and decision support tools. Where applied, these reduce...   | How should response time, decision speed, or implementation pace improve?<br>Improve decision speed and implementation pace to support to enable secure within required timelines. The contractor engineer supporting a government...   |

**Figure 2. Readiness Canvas for User Story # 1 for Level 1 Maturity**

**Summary Outputs.** Figure 3 depicts the operational features that include decisions, training recommendations, and implementation plans based on completed canvas data. It supports export to PDF, Word, and JSON formats for integration with existing acquisition documentation workflows.






**Figure 3. Operational Features—Decisions, Training, and Timeline**

This digital tool represents a key contribution of the research by translating the conceptual IDMP framework into a practitioner-ready instrument that acquisition professionals can use in portfolio reviews, strategy sessions, and workforce development activities.

### Six Dimensions of the Readiness Canvas

The IDMP Readiness Canvas organizes assessment across six dimensions, each addressing a critical facet of acquisition readiness. Together with four assessment elements (Current State, Issues, Capabilities, and Objectives), these dimensions form a 24-cell assessment matrix that provides comprehensive readiness coverage:

- **Capability (“Can we do it?”).** This dimension assesses the organization’s current ability to execute acquisition activities. It examines what can be done today, where capability fails or does not scale, what tools and processes enable execution, and what capabilities must be strengthened or matured. Capability assessment directly supports PAE portfolio evaluations by providing structured visibility into executable capacity across programs.
- **Evidence (“What proof exists?”).** This dimension evaluates the quality, availability, and traceability of evidence supporting acquisition decisions. It examines what evidence exists today, where evidence is weak or unreliable, what evidence pipelines and measures are available, and what evidence quality must improve. Evidence assessment supports the Acquisition Transformation Strategy’s emphasis on data-driven decision-making by ensuring that readiness claims are substantiated by verifiable data.
- **AI Role (“How does AI help?”).** This dimension captures the current and planned use of artificial intelligence within acquisition processes. It assesses what AI is currently being used, where AI misleads or lacks trust, what AI functions reliably augment decisions, and how AI support should evolve. This dimension is particularly relevant as the Acquisition Transformation Strategy emphasizes using AI to reduce administrative burden and accelerate decision-making (DoW, 2025a).
- **Human Role (“Where must humans decide?”).** This dimension examines how human judgment, authority, and oversight operate within acquisition environments. It evaluates how humans currently compensate and decide, where human authority is unclear or overloaded, what human expertise and review processes are formalized, and what human roles need clarification or strengthening. This dimension ensures that AI-augmented processes maintain appropriate human-in-the-loop governance, consistent with responsible AI principles in defense decision-support contexts (Hare et al., 2022; Preece et al., 2020).
- **Risk (“What could go wrong?”).** This dimension assesses the risk landscape across acquisition portfolios. It examines what risks are currently exposed, where risks are unmanaged or underestimated, what controls and mitigations are in place, and what risks should be anticipated earlier. The risk dimension incorporates considerations

aligned with the DoW Risk Management Framework (DoWI 8510.01) and Joint Risk Analysis standards, providing structured risk visibility for portfolio-level decision-making.

- **Time (“How fast?”).** This dimension evaluates the temporal aspects of acquisition decision-making and execution. It assesses how long response and decision-making currently take, where delays create harm or lost opportunity, what currently shortens cycle time, and how response time should improve. Time assessment directly supports the Acquisition Transformation Strategy’s emphasis on speed, addressing the critical need to deliver capabilities at the pace required by modern operational environments (DoW, 2025a).

Each cell in the 24-cell matrix includes a guiding question that varies by the selected maturity level (L1 through L5), an AI-generated draft response drawn from evidence sources, a human-validated response, a confidence score, source references with provenance tracking, identified gaps, and a validation status progressing from draft to validated. This per-cell structure ensures that every assessment is traceable, evidence-based, and subject to human oversight.




**Figure 4. IDMP Canvas: 24-Cell Assessment Matrix**

The canvas also generates three action components from the completed 24-cell analysis: Key Decisions (what must be decided, including owners, deadlines, and trade-offs), a Training Component (identifying human capability gaps and targeted uplift needs), and an Implementation Timeline (establishing when to act, with dependencies and scenarios).

### AI Processing Pipeline: From Context to Evidence

The IDMP User Story Canvas is supported by a multi-stage AI processing pipeline that transforms user-created inputs into structured, evidence-based readiness assessments. The pipeline architecture reflects the research team’s approach to integrating retrieval-augmented generation (RAG), orchestrated multi-actor processing, and continuous learning within a human-in-the-loop framework.

**Input Layer: Context Header.** The pipeline begins with human-created inputs captured through the canvas interface. These inputs include the mission-level and operational-level user stories, the selected perspective (Government, Prime Contractor, Vendor, Test/Evaluation, or Sustainment), the story trigger (Milestone, Risk Threshold, AI Anomaly, or Data Constraint), and the target maturity level. Together, these inputs define the problem context and scope the subsequent AI analysis.

**Context Processing.** The input layer feeds into a context processing stage that performs query reframing and context binding. This stage takes the user story, trigger, and maturity level as inputs and produces structured context and dimension-specific queries as outputs. Context processing ensures that downstream retrieval and analysis are precisely scoped to the user's operational environment and maturity target.

**Multiple RAG Models.** Structured queries are processed through multiple RAG models that retrieve relevant content from authoritative data sources. These sources include the DoW Digital Engineering Strategy, the DoWI 5000 Series, the Risk Management Framework (RMF), the Mission Engineering Guide v2.0, and industry benchmarks. Each source is indexed using FAISS (Facebook AI Similarity Search) vector stores to enable efficient semantic retrieval. The RAG models contextualize user stories by aligning them with prior cases, benchmarks, and established readiness guidance, producing evidence bundles that ground subsequent analysis in authoritative documentation.




**Figure 5. IDMP Processing Flow: Context to Evidence**

**Orchestrated Processing Functions.** Evidence bundles pass through three specialized processing functions operating as orchestrated multi-actor processors:

- **Standardization** harmonizes data from diverse DoW formats and sources to enable domain consistency, ensuring that readiness indicators from different programs and organizations can be meaningfully compared.
- **Tech Matching** correlates maturity gaps with appropriate technologies, tools, or training, and aligns maturity models with DoW methodologies to identify pathways for capability improvement.
- **Risk Assessment** performs AI-enabled risk scoring and prioritization, incorporating the RMF (DoWI 8510.01), Digital Engineering Strategy, Mission Engineering Guide v2.0, and Joint Risk Analysis standards to provide contextualized risk evaluation.

**Six Dimension Actors.** Processing function outputs are distributed to six specialized dimension actors, one for each readiness dimension (Capability, Evidence, AI Role, Human Role, Risk, and Time). Each actor generates draft responses for the four assessment elements within its dimension, producing the initial content for the 24-cell canvas matrix.

**Output Layer: Decision Support and Adaptive Insights.** The pipeline produces a comprehensive set of decision-support outputs including digital maturity insights, intelligent recommendations, predictive analytics and scenario forecasts, knowledge graph and reasoning

summaries, risk intelligence dashboards, technology alignment maps, IDMP Canvas updates, and HITL feedback channels. These outputs are designed to provide acquisition professionals with actionable, evidence-based decision support that integrates multiple analytical perspectives.

**Continuous Learning and Feedback Loop.** The architecture incorporates a continuous learning mechanism that feeds outcomes back into the system. This includes HITL validation results, model retraining and optimization data, user story evolution tracking, knowledge graph expansion, decision outcome tracking, and governance and oversight updates. This feedback loop ensures that the system improves over time as more user stories are processed and validated, creating an increasingly robust knowledge base for acquisition decision support.

## Discussion

The preliminary findings suggest that the IDMP framework addresses a meaningful gap in current acquisition decision-support capabilities. By combining structured user stories with integrated maturity assessments and AI-enabled analysis, the framework provides a mechanism for translating policy intent into operational practice.

Several aspects of the results merit further discussion:

- **Integration of Readiness Domains.** The IDMP framework's ability to synthesize readiness indicators across multiple domains—technology, manufacturing, integration, workforce, and organizational addresses a persistent limitation of existing readiness models. By structuring assessment within a unified maturity pathway, the framework provides portfolio leaders with a more comprehensive view of capability readiness than any single readiness model can provide independently (Sausser et al., 2009; See, 2021).
- **User Stories as Decision Artifacts.** The research demonstrates that user stories can function as more than requirements documentation. When structured within the IDMP framework and enriched through AI-enabled analysis, user stories become analytical artifacts that inform maturity assessment, capability gap analysis, and portfolio prioritization. This extends the traditional application of user stories beyond agile software development into complex defense acquisition environments (Sporsem et al., 2025).
- **AI-Enhanced Analysis with Human Oversight.** The AI/ML component's ability to identify patterns, cluster related stories, and assess readiness conditions demonstrates the potential for machine-assisted analysis in acquisition decision-making. Importantly, the human-in-the-loop approach ensures that AI-generated insights are validated by subject-matter experts, maintaining the trust and accountability necessary for high-stakes defense decisions (Hare et al., 2022; Preece et al., 2020).
- **Practitioner Readiness.** The IDMP User Story Canvas translates the conceptual framework into a practical tool that acquisition professionals can use within existing workflows. The structured 24-cell assessment matrix provides a repeatable methodology for readiness evaluation, while the AI processing pipeline automates analytical steps that would otherwise require significant manual effort.
- **AI-Enabled Analysis.** AI-enabled analysis has improved the consistency of readiness classification and the identification of cross-cutting themes across stakeholder groups. The combination of clustering, pattern recognition, and maturity scoring demonstrates the potential for AI to augment, but not replace, expert judgment in acquisition decision-support.





## Limitations and Future Research

As this research is ongoing with data collection continuing through September 2026, several preliminary limitations should be acknowledged.

The sample of workshop and survey participants, while diverse, does not represent the full breadth of the defense acquisition enterprise. Workshop scope is constrained to the specific acquisition transformation topics addressed in facilitated sessions and may not capture all relevant dimensions of stakeholder experience. The AI-enabled analysis is currently exploratory, and further validation is needed to assess the reliability and generalizability of AI-generated readiness classifications across broader acquisition contexts.

Future research should expand the stakeholder base to include small to medium businesses, additional acquisition organizations and operational commands, refine the IDMP maturity levels through broader empirical validation, and evaluate the framework's effectiveness in supporting actual portfolio-level acquisition decisions. The integration of additional AI capabilities including predictive analytics and natural language processing for automated user story quality assessment represents a promising direction for enhancing the framework's analytical power.

## Conclusion

This ongoing research demonstrates that the IDMP framework provides a mechanism for operationalizing acquisition transformation through structured user stories, integrated maturity assessment, and AI-enabled analysis. The framework addresses a persistent gap between policy objectives and practical decision-support capabilities by providing acquisition leaders with a repeatable methodology for assessing readiness across multiple domains.

Key contributions include: a validated approach for developing and analyzing acquisition user stories that capture operational perspectives on transformation challenges; an integrated maturity model that synthesizes readiness indicators across technology, manufacturing, integration, workforce, and organizational dimensions; evidence that AI-enabled analysis can enhance the consistency and depth of user story interpretation; and a practitioner-ready digital tool the IDMP User Story Canvas that translates the conceptual framework into an operational instrument.

With additional workshops, an AI/ML prototype demonstration, and a stakeholder survey planned for the remainder of the performance period, the research team anticipates further refinement and validation of these preliminary findings.

As the DoW continues to implement its acquisition transformation strategy, frameworks such as the IDMP provide a structured approach for ensuring that reform objectives are translated into measurable readiness improvements and actionable decision support. By centering the acquisition process on stakeholder perspectives and integrating multiple readiness dimensions within a unified analytical framework, the IDMP framework contributes to ongoing efforts to accelerate the delivery of capabilities that support national defense objectives.

## References

- Amna, A. R., Poels, G., & Quispe-Otazu, R. (2022). Systematic literature mapping of user story research. *IEEE Access*, 10, 75287–75302.  
<https://doi.org/10.1109/ACCESS.2022.3191649>
- Cohn, M. (2004). *User stories applied: For agile software development*. Addison-Wesley.



- Dahmann, J. S., & Parasidis, G. I. (2024). Mission engineering: A technical perspective. *The ITEA Journal of Test and Evaluation*, 45(3). <https://itea.org/journals/volume-45-3/mission-engineering/>
- DoW. (2020). *DoW Instruction 5000.88: Engineering of defense systems*. <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500088p.PDF>
- DoW. (2023). *Mission engineering guide* (version 2.0). [https://ac.cto.mil/wp-content/uploads/2023/11/MEG\\_2\\_Oct2023.pdf](https://ac.cto.mil/wp-content/uploads/2023/11/MEG_2_Oct2023.pdf)
- DoW. (2025a). *Acquisition transformation strategy*. <https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/ACQUISITION-TRANSFORMATION-STRATEGY.PDF>
- DoW. (2025b). *Transforming the defense acquisition system into the warfighting acquisition system*. <https://media.defense.gov/2025/Nov/10/2003819439/-1/-1/1/TRANSFORMING-THE-DEFENSE-ACQUISITION-SYSTEM-INTO-THE-WARFIGHTING-ACQUISITION-SYSTEM-TO-ACCELERATE-FIELDING-OF-URGENTLY-NEEDED-CAPABILITIES-TO-OUR-WARRIORS.PDF>
- Department of War, Office of the Chief Technology Officer. (2025). *Technology readiness assessment guidebook*. <https://www.cto.mil/wp-content/uploads/2025/03/TRA-Guide-Feb2025.v2-Cleared.pdf>
- Hallmann, D., Jacob, K., Lüttgen, G., Schmid, U., & von der Weth, R. (2025). USeR: A web-based user story eReviewer for assisted quality optimizations. In *Proceedings of the 10th International Conference on Model-Driven Engineering and Software Development*. <https://doi.org/10.5220/0012619200003693>
- Hare, J. Z., Kase, S. E., Rinderspacher, B., Hung, C., Su, S., & Krayzman, T. (2022). The future of collaborative human-artificial intelligence decision-making for mission planning. *Frontiers in Psychology*, 13, 850628. <https://doi.org/10.3389/fpsyg.2022.850628>
- Joint Defense Manufacturing Technology Panel. (2025). *Manufacturing readiness level deskbook*. [https://www.dodmrl.com/MRL\\_Deskbook\\_2025.pdf](https://www.dodmrl.com/MRL_Deskbook_2025.pdf)
- Kasauli, R., Knauss, E., Horkoff, J., Liebel, G., & Neto, F. D. O. (2021). Requirements engineering challenges and practices in large-scale agile system development. *Journal of Systems and Software*, 172, 110851. <https://doi.org/10.1016/j.jss.2020.110851>
- Kuhail, M. A., & Lauesen, S. (2022). User story quality in practice: A case study. *Software*, 1(3), 215–233. <https://doi.org/10.3390/software1030010>
- Lavin, A., Gilligan-Lee, C. M., Visnjic, A., Ganju, S., Newman, D., Baydin, A. G., Gal, Y., & others. (2022). Technology readiness levels for machine learning systems. *Nature Communications*, 13, 6039. <https://doi.org/10.1038/s41467-022-33128-9>
- Paulo, E., & Beery, P. T. (2019). Application of model-based systems engineering concepts to support mission engineering. *Systems*, 7(3), 44. <https://doi.org/10.3390/systems7030044>
- Preece, A., Braines, D., Cerutti, F., Pearson, G., Kaplan, L. M., Srivastava, M., Tomsett, R. J., & others. (2020). Rapid trust calibration through interpretable and uncertainty-aware AI. *Patterns*, 1(4), 100049. <https://doi.org/10.1016/j.patter.2020.100049>
- Qian, L., Xu, X., Dou, Y., Tan, Y., Ma, Y., & Zhang, Z. (2023). A requirement quality assessment method based on user stories. *Electronics*, 12(10), 2155. <https://doi.org/10.3390/electronics12102155>
- Sauser, B. J., Ramirez-Marquez, J. E., Magnaye, R. B., & Tan, W. (2009). Defining an integration readiness level for defense acquisition. *Information Knowledge Systems Management*, 8, 101–114.



- See, J. E. (2010). *The development and initial evaluation of the human readiness level scale* (Doctoral dissertation, Naval Postgraduate School). Calhoun.
- See, J. E. (2021). *Human readiness levels explained*. Sandia National Laboratories.  
<https://www.osti.gov/servlets/purl/178752>
- Systems Engineering Research Center (SERC). (2023). *Systems engineering modernization: Acquisition with digital engineering (WRT-1058 and WRT-1057.18g)*.  
[https://sercuarc.org/wp-content/uploads/2025/10/T1\\_4.15pm\\_SRR-2023-SERC-McDermott-1057\\_18-1058-SE-Modernization.pdf](https://sercuarc.org/wp-content/uploads/2025/10/T1_4.15pm_SRR-2023-SERC-McDermott-1057_18-1058-SE-Modernization.pdf)
- Sporsem, T., Karlsen, T. J., & Moe, N. B. (2025). User stories as boundary objects in agile requirements engineering. *Journal of Systems and Software*, 219, 112107.
- The White House. (2025, April 9). *Modernizing defense acquisitions and spurring innovation in the defense industrial base*. <https://www.whitehouse.gov/presidential-actions/2025/04/modernizing-defense-acquisitions-and-spurring-innovation-in-the-defense-industrial-base/>



# Augmented Acquisition: Operationalizing Human-in-the-Loop AI to Accelerate Warfighting Capabilities

**Omar Haroun**—is the CEO and Co-Founder of Eudia, a company building augmented intelligence to empower humans at the world's most important organizations. After the successful exit of his previous company, Text IQ (acquired by Relativity in 2021), Omar founded Eudia to transform how legal and regulatory work is performed by positioning AI as a tool that supports human judgment to deliver maximum productivity. Omar is a visionary whose life work has been centered on building great technology that contributes to societal advantage. Prior to Text IQ, he served as Chief Revenue Officer at Greatist and co-founded Sportaneous. He has been recognized as a White House Champion of Change and holds a JD from Columbia Law School, an MBA from Columbia Business School, and an MA and BA in Philosophy, Politics, and Economics from the University of Oxford.

**Ashish Agrawal**—is the Co-Founder and CTO of Eudia, where he leads the technical vision and development of the company's augmented intelligence platform. Ashish is an AI and technology leader with two decades of experience across Google, Amazon, Apple, IBM Research, and various early-stage AI companies. At Amazon, he led the technical direction for Alexa's AI capabilities, including conversation, speech, and natural language understanding. At Apple, he built highly secure local search functionality for Maps, Siri, and Safari, serving hundreds of millions of users. At Cresta, he directed a team building real-time generative AI-driven knowledge assistance. Earlier in his career, Ashish served as an early engineer on Google Finance. He holds a master's in computer science from ETH Zurich and a bachelor's with honors from the Indian Institute of Information Technology, Allahabad. He has numerous patents and publications in AI and machine learning.

**Dylan Mrla**—is a Product Manager at Eudia and a United States Air Force veteran. He drives product strategy and development across Eudia's platform, with direct ownership of the company's defense acquisition vertical. Dylan has emerged as a transformational leader in applying emerging AI technology to highly regulated workflows, drawing on his firsthand experience as a warranted Contracting Officer to deeply understand how to augment human talent. During his active-duty service, Dylan served as Deputy Director of the Operational Contract Support Integration Cell for Combined Joint Task Force Operation Inherent Resolve, coordinating contracting operations across Kuwait, Iraq, Syria, and Jordan in direct support of combat operations. Stateside, he oversaw complex test and evaluation contracting. His experience at the strategic, enterprise, and operational levels of the acquisition process were critical for this research. His unique combination of domain knowledge and AI product expertise positions him at the forefront of efforts to modernize high-level professionals through human-centered technology, particularly in defense acquisition. Mrla holds a BA in Management from Michigan State University's Eli Broad College of Business, where he was recognized by his undergraduate student body as the honored commencement speaker.

## Abstract

Contracting officers are the backbone of delivering mission-critical advantages to the warfighter, responsible for navigating thousands of pages of regulatory guidance across the Federal Acquisition Regulation (FAR) and its sub-regulatory supplements while producing accurate, defensible, and timely acquisition packages. Beyond regulatory complexity, contracting professionals must manage massive volumes of contract data across the enterprise and at the squadron level while often simultaneously operating across multiple disconnected systems to accomplish a single task. The acquisition process itself spans numerous stages, from requirements development through pre-solicitation, solicitation, pre-award, and award, each with distinct documentation and compliance demands.

Compounding these challenges, requirement owners frequently delay timelines of critical acquisitions due to poorly defined or incomplete requirements, a key area where AI can provide meaningful augmentation. As the Department of War accelerates efforts to deliver warfighting capability at the speed of relevance, the acquisition workforce faces mounting pressure to reduce Procurement Acquisition Lead Time (PALT) without sacrificing compliance or accountability. The



leading approach to address these challenges with cutting-edge AI technology is to combine human input with artificial intelligence to significantly enhance the productivity of high-performing personnel.

This paper investigates an alternative paradigm: human-centered AI that augments contracting officer judgment rather than replacing it. The proposed approach positions AI as a copilot that reviews, validates, and enhances human-generated content in real time, enabling acquisition professionals to reach their maximum capacity for production while retaining full ownership of every decision. Our research is grounded in an active research and development (R&D) contract with the 39th Contracting Squadron (39 CONS) at Incirlik Air Base, where Eudia is developing and evaluating an augmented acquisition capability comprising three integrated applications: Insights, Sigma, and Augmented Contract Review (ACR). The system reviews human-generated acquisition documents against structured training data drawn from the FAR, DFARS, DAFFARS, internal policy, and historical contract files, providing clear citations, plain language explanations of regulatory risks, and structured feedback that strengthens document quality and compliance confidence.

Through testing, we estimate PALT reductions varying by acquisition type—with simpler actions such as Simplified Acquisition Procedures under \$250,000 showing different improvement rates than complex Source Selection efforts. Preliminary findings indicate that human-in-the-loop AI reduces document error rates, accelerates regulatory review, and supports workforce development by reinforcing training concepts and surfacing relevant guidance for junior personnel. This research contributes to the symposium's focus on accelerating warfighting capability by demonstrating that the fastest path to contract award is not replacing human judgment but amplifying it.

**Keywords:** Defense acquisition, human-centered AI, augmented intelligence, contracting officer decision support, FAR compliance, DFARS compliance, DAFFARS compliance, Procurement Acquisition Lead Time, PALT reduction, human-in-the-loop AI, acquisition workforce, regulatory compliance, warfighting capability, 39th Contracting Squadron, document validation, AI copilot, acquisition acceleration.

## Introduction

The Department of War acquisition enterprise faces an enduring imperative: deliver warfighting capabilities to the field faster than adversaries can adapt. Yet the contracting process governed by the Federal Acquisition Regulation (FAR) and its sub-regulatory supplements remains one of the most labor-intensive functions in the defense enterprise. While the defense budget and volume of acquisition requirements have expanded significantly, the size of the acquisition workforce has not grown proportionally. Contracting officers must navigate thousands of pages of regulatory guidance, produce extensive documentation across multiple acquisition phases, and ensure every action can withstand audit scrutiny and legal challenge. This work is performed while managing massive volumes of historical contract data, operating across disconnected information systems, and coordinating with requirement owners whose input quality directly impacts acquisition timelines.

The manual nature of acquisition documentation creates significant bottlenecks. Contracting professionals currently cross-check documents against regulatory checklists through tedious, line-by-line review. They search for precedent across filing systems that may contain years of institutional knowledge but lack intuitive retrieval mechanisms. They validate clause selection, citation accuracy, and justification language through processes that are both time-consuming and susceptible to human error. As the Department of War accelerates its push toward rapid capability deployment, these people-intensive compliance workflows struggle to keep pace, introducing delays that directly impact mission readiness.

The challenge extends beyond the contracting office. Defining requirements is challenging, particularly in higher-intensity mission environments. When requirement owners





struggle to articulate their needs, contracting officers must interject during the requirements development stage which can add days or weeks to Procurement Acquisition Lead Time (PALT). Pre-solicitation activities and documentation each present opportunities for error or delay. Across every stage of the acquisition life cycle—requirements development, pre-solicitation, solicitation, pre-award, and award—a significant gap exists between the speed at which the enterprise must operate and the capacity of the current workforce to produce compliant, defensible documentation while achieving contract award timelines comparable to commercial industry.

Artificial intelligence presents an opportunity to close this gap, but how AI is implemented matters profoundly. Approaches that seek to fully automate document generation such as producing market research reports, sole source justifications, or pricing memorandums with minimal human involvement other than prompting introduce risks that may ultimately slow rather than accelerate acquisition. When AI generates entire template documents, questions arise about explainability, auditability, and the statutory accountability that warranted contracting officers bear for every action they sign. The most promising path forward is not automation that replaces human judgment but augmentation that amplifies it.

## Research Context and Scope

This paper presents findings from an active research and development effort to operationalize human-in-the-loop AI within defense acquisition. Eudia is a Silicon Valley-based technology company founded in 2023 with the mission of augmenting humans at the world's most important organizations to achieve their full potential. In 2025, Eudia was selected for a Direct-to-Phase II Small Business Innovation Research (SBIR) contract by the United States Air Force to develop AI-augmented acquisition technology for the 39th Contracting Squadron (39 CONS) at Incirlik Air Base in Türkiye. Eudia's selection underscores the company's track record of working hands-on with elite organizations to build the next generation of high-productivity, human-centered workforces.

The 39 CONS engagement serves as the primary operational testing environment for this research. As an operational contracting squadron, 39 CONS executes a diverse portfolio of contract actions supporting base operations, providing an authentic environment where AI tools must perform under real-world conditions with actual mission consequences. The research team embedded directly with contracting personnel to understand workflow pain points and document quality challenges and opportunities for AI augmentation. This collaboration culminated in the successful deployment of multiple operational prototypes that validated a core hypothesis: human augmentation represents the future of AI in defense acquisition.

To ensure the research captured perspectives beyond operational contracting, the team conducted extensive discovery sessions with mission partners across the acquisition enterprise:

- Space Systems Command (SSC): Examining utility in systems acquisition and complex program management environments
- Office of the Chief of Space Force Contracting: Understanding service-level contracting policy and strategic priorities
- Air Force Test Center (AFTC): Assessing application to test and evaluation contracting with specialized technical requirements
- Headquarters Air Force Materiel Command (HQAFMC/PK): Exploring major command contracting policy and oversight considerations



- Air Force Life Cycle Management Center Armament Directorate (AFLCMC/EB): Aligning with program-level acquisition strategy and policy direction

These engagements validated that the challenges observed at 39 CONS—regulatory complexity, disconnected systems, requirement quality issues, and the need for faster compliance validation—exist across the acquisition enterprise. They also confirmed that human-centered AI, properly implemented, could address these challenges in contexts ranging from operational contracting to systems acquisition and Other Transaction Authority (OTA) contracting motions. Beyond the Department of War, the research team met with the offices of 13 members of Congress, where there was shared recognition that the federal acquisition process is a critical problem that is ripe for innovation. The consistency of this perspective from flight-level contracting offices to the halls of Congress demonstrates that acquisition modernization is a systemic imperative. This broad alignment creates a compelling case for human-centered AI as a solution that meets both operational needs and strategic policy objectives.

## Research Objectives

This research investigates how AI systems designed to amplify the throughput of a human asset can meaningfully accelerate defense acquisition while maintaining the compliance, transparency, and accountability the process requires. The study examines Eudia's Unified Platform, which comprises three integrated applications:

- **Contract Insights:** enables natural language queries against historical contract and market research data, displays relevant contract information in an organized tabular view that automatically refreshes, and creates a contracts dashboard filterable by DODACC.
- **Projects:** supports requirement owners in producing higher-quality inputs to the contracting process and allows end-users to query against historical requirement and contract data.
- **Augmented Contract Review (ACR):** provides real-time compliance validation and risk identification for human-generated acquisition documents.

The sections that follow present the regulatory and operational context for this research, describe the methodology employed, detail the technical implementation of the Eudia platform, and report preliminary findings on PALT reduction, document quality improvement, and workforce development benefits. The paper concludes with implications for broader adoption of human-centered AI across the defense acquisition enterprise.

Figure 1 illustrates the Sigma application interface, designed to support requirement owners in developing higher-quality inputs to the acquisition process. In defense acquisition, the quality of requirements directly impacts every downstream activity from market research and solicitation development to source selection and contract award. When requirement owners struggle to articulate clear, complete, and compliant requirements, contracting officers must intervene to clarify intent, identify missing information, and reconcile inconsistencies. These iterations add weeks or months to Procurement Acquisition Lead Time before the formal contracting process even begins and often during.

It is important to note that Sigma and the broader Eudia platform are designed to feed insights into the human review process, not replace it. The goal is to improve the effectiveness of both requirement owners and contracting professionals while allowing human experts to focus on judgment calls, strategic decisions, and mission alignment. In short, augmenting human effort through AI reduces risk in acquisition programs by ensuring that potential issues



are identified early, best sources are determined with maximum efficiency, and the documentation gauntlet is executed quickly, all while preserving human authority over every decision point. Delays and compliance issues can surface at many stages of the acquisition life cycle. Department of War policies drive requirements but also shape acquisition strategy, market research, solicitation structure, evaluation criteria, and contract administration. When regulations require specific processes, clauses, or documentation standards, these must be accurately reflected across all relevant acquisition documents. The complexity of this regulatory environment combined with the volume of actions a contracting squadron must execute creates significant opportunity for AI augmentation.




**Figure 1: Eudia's Sigma Projects User Interface**

Table 1 presents baseline Procurement Acquisition Lead Time data provided by the Air Force Test Center (AFTC). These figures represent target timelines for various acquisition types and illustrate the scope of the challenge. Even modest percentage reductions in these lead times, when applied across hundreds of contract actions, translate directly into faster delivery of warfighting capability at lower cost to the taxpayer. Early estimates suggest that integrating the Eudia Unified Platform at scale could reduce average operational contracting squadron PALT by approximately 31%, and that is just with initial roll out. The platform gets smarter through usage.

**Table 1: Average PALT Across AFTC**

| Acquisition Type                                                    | Lead Time (Days) |
|---------------------------------------------------------------------|------------------|
| Simplified Acquisition Procedures (SAP) (FAR Part 13 < \$350K)      | 30               |
| Simplified Acquisition Procedures (SAP) (FAR subpart 13.5 > \$350K) | 180              |
| Contracting by Negotiation (Non-Competitive) (FAR Part 15)          | 180              |
| Source Selection (FAR Part 15 Competitive)                          | 240              |
| Broad Agency Announcement (BAA)                                     | 270              |
| General Service Administration (GSA) (FAR Part 8)                   | 30               |
| Small Business Innovative Research (SBIR) Phase II                  | 90               |
| Delivery Orders (DOs) – Non-Negotiated                              | 30               |
| Task Orders (TOs) – Firm Fixed Price/Cost                           | 75/180           |
| Change Order (13A, SF30) – Non-commercial                           | 60               |



|                                           |     |
|-------------------------------------------|-----|
| Administration (13B, SF30)                | 10  |
| Supplemental: Time Extension (13C, SF30)  | 30  |
| Supplemental: Change to Value (13C, SF30) | 120 |
| Supplemental: All Others (13C, SF30)      | 30  |
| Other: Incremental Funding (13D, SF30)    | 10  |
| Other: All Other Authority (13D, SF30)    | 10  |

## Driving Human-Centered AI into the Acquisition Life Cycle

Human-centered AI represents a fundamental shift in how artificial intelligence is deployed within high-stakes professional environments. Rather than positioning AI as an autonomous decision-maker that generates outputs for human approval, human-centered AI positions the human as the primary actor and AI as a force multiplier that enhances human capability. This is achieved through structured training and an iterative development approach in direct partnership with operational users. In the context of defense acquisition, this distinction carries significant implications for compliance, accountability, and operational effectiveness.

Three core principles define human-centered AI in acquisition:

- **Augmentation over Automation.** The fundamental mindset that AI is best utilized to enhance human productivity rather than replacing human judgment. For example, rather than generating a sole source justification from scratch, the AI reviews a human-drafted justification and identifies missing regulatory citations, incomplete market research references, or language that may not withstand audit scrutiny. The contracting officer & requirement owner retain ownership of the document while benefiting from comprehensive compliance validation that would otherwise require hours of manual cross-referencing.
- **Explainability and Transparency.** Every recommendation or finding produced by Eudia includes clear citations to regulatory authority and plain-language explanations of why something may constitute a risk. When the system flags a potential issue such as a missing clause or an inconsistency between the statement of work and the evaluation criteria, the contracting officer can immediately understand the basis for the finding and make an informed decision about how to proceed. This transparency is essential for maintaining audit trails and defending contracting decisions.
- **Human Authority and Accountability.** Eudia advises but never decides. Warranted contracting officers hold statutory authority to obligate government funds, and this authority cannot be delegated to a machine. Human-centered AI preserves the contracting officer's role as the final decision-maker while providing them with capabilities that exceed what any individual could accomplish alone. Every recommendation can be accepted, modified, or rejected based on the contracting officer's professional judgment.

These principles are essential in an environment where the consequences of error are significant. A missing clause can result in a bid protest. An incomplete justification can trigger audit findings. A poorly defined requirement can delay contract award by weeks or months. Human-centered AI addresses these risks not by removing humans from the process but by ensuring that humans operate at their maximum capacity.



**Table 2: How Human-Centered AI Maps to Key Stages of the Acquisition Life Cycle**

| Acquisition Stage        | Human Activity                                                           | AI Augmentation Value                                                                                           |
|--------------------------|--------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Requirements Development | Requirement owner defines the need                                       | Sigma validates completeness, identifies ambiguities, and flags language that may cause downstream issues       |
| Market Research          | Contracting officer researches potential sources                         | Insights surfaces historical precedent, similar acquisitions, and relevant market data from past contract files |
| Pre-Solicitation         | Contracting officer prepares solicitation documents                      | ACR validates clause selection, citation accuracy, and regulatory alignment in real time                        |
| Solicitation             | Contracting officer issues and manages solicitation                      | Insights enables rapid response to industry questions using historical Q&A patterns                             |
| Evaluation & Pre-Award   | Contracting officer evaluates proposals and prepares award documentation | ACR reviews pricing memoranda, source selection documentation, and justifications for compliance risks          |
| Award                    | Contracting officer executes award                                       | ACR provides final compliance validation before signature                                                       |

### Exploring Human-Centered AI for Acquisition Compliance

Exploring specific examples provides important insights into the opportunities and current capabilities of human-centered AI in defense acquisition regulatory compliance. Through our Direct-to-Phase II SBIR contract with the 39th Contracting Squadron, we have developed and tested the Eudia Unified Platform in an operational environment, processing real acquisition documents against the FAR and its sub-regulatory supplements. Although testing and refinement continue, our initial results demonstrate significant potential for AI to augment contracting professionals across the acquisition life cycle.

The Eudia Unified Platform comprises three integrated applications, each designed to address specific pain points identified through direct collaboration with operational users:

Insights enables contracting officers to query historical contract data and regulatory guidance using natural language. Rather than manually searching through years of contract files, shared drives, and policy memoranda, contracting personnel can ask targeted questions and receive contextually relevant answers with citations to source documents. This capability transforms institutional knowledge previously locked in filing systems or the memories of career professionals into an accessible, searchable resource. Sigma supports requirement owners in developing complete, compliant, and clearly articulated requirements before they enter the formal contracting process. By identifying ambiguities, missing information, and language likely to cause downstream issues, Sigma reduces the iterations between requirement owners and contracting officers that frequently delay acquisition timelines. Augmented Contract Review (ACR) provides real-time compliance validation for human-generated acquisition documents. As contracting officers draft market research reports, sole source justifications, pricing memoranda, and other deliverables, ACR reviews the content against the FAR, DFARS, DAFFARS, local policy, and historical precedent. The system identifies missing clauses, citation errors, incomplete justification language, and potential compliance risks each finding accompanied by regulatory citations and plain-language explanations.

A common use case illustrates the value of this approach. When a requirements owner prepares a sole source justification under FAR Part 6, they must document specific statutory authorities, demonstrate that requirements are not unduly restrictive, and provide evidence that





market research was conducted. Traditionally, this requires cross-referencing multiple regulatory sections, reviewing historical justifications for similar actions, and manually validating that all required elements are present. ACR performs this validation in real time, flagging gaps or risks while the document is being developed rather than during peer review or legal review when corrections are more costly and time-consuming.

**Table 3. Examples of Compliance Issues Identified by the Eudia Platform Across Different Acquisition Document Types**

| Document Type             | Issue Identified                 | AI Finding and Explanation                                                                                                                                                                                             |
|---------------------------|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Sole Source Justification | Missing market research citation | The justification references market research but does not cite the specific report or date. FAR 6.303-2 requires documentation of market research supporting the sole source determination.                            |
| Market Research Report    | Incomplete source consideration  | The report identifies three potential sources but does not document why two were determined non-responsive. FAR Part 10 requires documentation of sources considered and rationale for narrowing the competitive pool. |
| Award Documentation       | Missing required clause          | The contract does not include FAR 52.204-25 (Prohibition on Contracting for Certain Telecommunications). This clause is required for all contracts exceeding the micro-purchase threshold.                             |

## Enterprise Portfolio Intelligence: A Critical Capability Gap

Through collaborative discussions with stakeholders at the Air Force Test Center and Space Systems Command, and Air Force Materiel Command, a critical insight emerged: contracting squadrons are among the most data-rich organizations in the federal government, yet they have almost no clear ways to leverage that data usefully. Every contract action generates documentation: requirements, market research, solicitations, proposals, evaluations, awards, modifications, and closeout files. Across an enterprise, this represents millions of data points containing institutional knowledge about pricing, performance, sources, compliance patterns, and risk indicators. Yet in practice, this wealth of information remains largely inaccessible, trapped in disconnected filing systems and individual memories.

The pain points identified across stakeholder engagements were remarkably consistent. Leaders described misaligned expectations between contracting offices and requirement owners, resulting in rework and delays. They noted the difficulty of tracking obligation rates and financial risk such as contracts requiring incremental funding across a portfolio of hundreds of active actions. They spoke candidly about overworked and undermanned contracting shops being asked to do more with less, without tools that provide meaningful visibility into workload distribution, bottlenecks, or emerging risks. And they acknowledged that despite years of investment in contracting information systems, squadron leadership often lacks a real-time, consolidated view of portfolio health.



These findings were not derived from policy documents or external reports but from direct conversations with acquisition professionals across the enterprise. From operational squadrons to major command headquarters, stakeholders articulated the same gap: contracting organizations are drowning in data but starving for insight. What leaders need is not another system that generates reports but a capability that transforms raw data into actionable intelligence with real-time visibility into portfolio health, proactive risk identification, and decision-quality information available on demand. The Insights dashboard within the Eudia Unified Platform was designed to address this exact need.

The research confirmed that what acquisition leaders need is not more data but more intelligence: the ability to see across their portfolio in real time, identify risks before they become problems, and make informed decisions about resource allocation and priorities. This is precisely the capability that the Insights dashboard within the Eudia Unified Platform was designed to deliver.

## Repository Awareness Through Natural Language Query

A foundational challenge in acquisition is accessing institutional knowledge. When a contracting officer begins work on a new requirement, relevant precedent almost certainly exists somewhere: a similar contract awarded 2 years ago, a justification that addressed comparable circumstances, a set of clauses used for a parallel requirement. But finding that precedent requires knowing where to look, having time to search, and recognizing relevance when it appears. In practice, contracting officers often start from scratch because locating historical examples takes longer than simply recreating the work.

The Insights AskAI capability addresses this challenge by enabling natural language queries against a squadron's complete contract repository. Rather than constructing Boolean searches or navigating folder hierarchies, contracting personnel can ask questions in plain language: "What justification language did we use for sole source IT services last year?" or "Which contracts with this vendor have had performance issues?" or "Show me examples of market research reports for similar supplies." The system retrieves contextually relevant results with citations to source documents, transforming years of institutional knowledge into an accessible resource. This capability has immediate practical value. Junior contracting officers gain access to the collective experience of the organization, accelerating their development and reducing reliance on senior personnel for routine guidance. Experienced officers save time on research, freeing capacity for higher-value activities. Peer reviewers can quickly validate whether a document aligns with organizational precedent. And leadership can query the repository to understand patterns: how the FSR support for LPWS systems were contracted last time it was competed, which requirement types consistently experience delays, or where compliance findings have historically concentrated.




**Figure 2: Insights AskAI Interface – Natural Language Query Capability**



## Portfolio Oversight Through Real-Time Dashboard

Beyond document-level queries, acquisition leaders require portfolio-level visibility. The Insights dashboard provides a consolidated, real-time view of contracting operations, displaying key metrics that enable informed decision-making and proactive risk management. As illustrated in Figure 3, the dashboard presents critical portfolio indicators including total obligations, contract ceiling values, active indefinite delivery vehicles and at-risk contract identification. Additionally, the dashboard is filterable by DODACC, customizable, and deployable with live data. Leadership can immediately see the distribution of spending by unit, identify top industry partners by obligation volume, and understand concentration risk across the vendor base. The system tracks unique contracts, categorizes actions by product service code (PSC), and highlights contracts flagged for risk such as those approaching ceiling limits, requiring incremental funding, or exhibiting performance concerns.

This portfolio-level intelligence addresses a gap that stakeholders consistently identified: the inability to see the forest for the trees. Individual contracting officers manage their assigned actions, but squadron leadership often lacks a mechanism to aggregate that information into actionable intelligence. The dashboard transforms raw contracting data into visual insights that support resource allocation, risk mitigation, and strategic planning. Critically, this capability exists today within the Eudia Unified Platform. It is not a future development or conceptual recommendation but an operational tool validated through the research engagement.




**Figure 3 Insight's Portfolio-Level Intelligence Dashboard**

## Implications for Acquisition Workforce Productivity

The combination of document-level query capability and portfolio-level dashboard intelligence represents a significant advancement in how contracting organizations can operate. Rather than spending hours searching for precedent, contracting officers can retrieve relevant examples in seconds. Rather than manually compiling status reports from disparate sources, leadership can access real-time portfolio views on demand. Rather than discovering risks reactively when a contract approaches its ceiling or a vendor underperforms, the system surfaces potential issues proactively.



This is the essence of human-centered AI: not replacing the judgment of acquisition professionals but providing them with capabilities that allow them to operate at maximum productivity. The contracting officer still makes every decision. The flight chief still sets priorities. The squadron commander still owns risk. But each operates with better information, faster access, and clearer visibility than previously possible. The research indicates that these capabilities, deployed at scale, have potential to fundamentally change acquisition workforce productivity. When contracting professionals spend less time searching and more time analyzing, less time compiling and more time deciding, the entire acquisition enterprise moves faster. Requirements reach contract sooner. Warfighters receive capability earlier. And the acquisition workforce, which is often cited as overworked and undermanned, gains capacity without adding headcount.

## Research Results

By April 2026, this research reflects approximately 5 months of operational data, user feedback, and system performance metrics gathered through the ongoing R&D contract with the 39th Contracting Squadron at Incirlik Air Base. While full multi-month impact analyses will continue over the period of performance, the early results demonstrate clear evidence that human-centered AI improves both the speed and quality of contracting workflows without replacing contracting officer judgment or statutory accountability.

The following sections present findings organized around the four research questions that guided this study.

**Research Question 1:** What regulatory, audit, and risk management gaps emerge when acquisition workflows move toward full automation?

Through interviews with acquisition professionals across the enterprise, from operational contracting squadrons to major command headquarters, this research identified significant concerns about automation-first approaches to AI in acquisition.

The most frequently cited gap involves explainability. When AI systems generate acquisition entire documents, contracting officers cannot fully explain the rationale behind the language produced. This creates vulnerability during audits, bid protests, and legal reviews, where contracting officers must defend not only what decision was made but how and why it was made. Automation-first approaches shift the burden of explanation from the document creator to a system that cannot be questioned or held accountable. A second gap involves statutory accountability. Warranted contracting officers hold legal authority to obligate government funds, an authority that cannot be delegated. When AI generates documents for human signature, the contracting officer's role shifts from decision-maker to reviewer. This inversion may satisfy procedural requirements but undermines the substantive judgment that warrants are intended to ensure. Multiple stakeholders expressed concern that automation-first approaches could erode professional expertise over time as contracting officers become accustomed to accepting AI outputs without critical evaluation.

A third gap involves auditability and documentation integrity. Acquisition documentation must demonstrate a clear chain of reasoning from requirement to award. AI-generated documents may produce plausible language without the evidentiary foundation that auditors require. When findings emerge, contracting officers may struggle to explain AI-generated content they did not author, creating compliance risk rather than reducing it. These gaps do not suggest that AI has no role in acquisition. Rather, they indicate that how AI is implemented matters profoundly. Human-centered AI designed to augment rather than replace addresses these gaps by keeping the contracting officer in the decision-making seat while providing capabilities that enhance their effectiveness.



**Research Question 2:** How can an AI-augmented, human-in-the-loop approach be structured to provide real-time review, risk insight, and compliance validation within contracting workflows?

This research developed and validated a three-application structure that addresses distinct needs across the acquisition life cycle while maintaining human authority at every decision point.

Sigma addresses the upstream challenge of requirement quality. By providing AI-augmented support to requirement owners as they develop statements of need, the system reduces ambiguity and incompleteness before requirements enter the formal contracting process. This intervention point is critical: delays caused by poorly defined requirements cannot be recovered downstream, making early augmentation essential to PALT reduction. Augmented Contract Review (ACR) provides real-time compliance validation as contracting officers develop acquisition documents. Rather than generating documents for human review, ACR reviews human-generated documents and provides structured feedback including regulatory citations, risk explanations, and recommended corrections. The contracting officer retains full ownership of the document while benefiting from comprehensive validation that would otherwise require hours of manual cross-referencing. Figure 4 illustrates the ACR interface, showing how findings are presented with clear citations and actionable guidance. The specific review shown in Figure 4 was done on a contract award. In second ACR generated a document-based risk assessment and flagged missed clauses, out of scope clauses included in the contract, and incorrect fill-in language, and made recommendations for changes to the CLIN structure. Each risk is shown in the side view and allows the user to easily click and see the source of the citation that is driving a flagged risk. This enables users to upskill through usage.

← 03b. Contract FA568525P0032 - Signed.pdf

Citation text not found in original document. Switch to processed view to see the highlighted text.

Original Processed Download

Page 11 of 21

Special Contract Requirements

Contract Clauses

FAR Clauses Incorporated by Reference

| Number    | Title                                                                     | Effective Date | Alternate/ Deviation | Variation Effective Date |
|-----------|---------------------------------------------------------------------------|----------------|----------------------|--------------------------|
| 52.204.13 | System for Award Management Maintenance                                   | Oct 2018       |                      |                          |
| 52.204.18 | Commercial and Government Entity Code Maintenance                         | Aug 2020       |                      |                          |
| 52.212.4  | Contract Terms and Conditions Commercial Products and Commercial Services | Nov 2023       |                      |                          |
| 52.223.5  | Pollution Prevention and Right-to-Know Information                        | May 2024       |                      |                          |

DFARS Clauses Incorporated by Reference

| Number       | Title                                                                                      | Effective Date | Alternate/ Deviation | Variation Effective Date |
|--------------|--------------------------------------------------------------------------------------------|----------------|----------------------|--------------------------|
| 252.203.7000 | Requirements Relating to Compensation of Former DoD Officials                              | Aug 2011       |                      |                          |
| 252.203.7002 | Requirement to Inform Employees of Whistleblower Rights                                    | Dec 2022       |                      |                          |
| 252.204.7004 | Anticorruption Business Training for Contractors                                           | Jan 2023       |                      |                          |
| 252.204.7012 | Safeguarding Covered Defense Information and Cyber Incident Reporting                      | May 2024       |                      |                          |
| 252.204.7016 | Prohibition on the Acquisition of Covered Defense Telecommunications Equipment or Services | Jan 2023       |                      |                          |
| 252.204.7020 | NIST SP 800-171 DoD Assessment Requirements                                                | Nov 2023       |                      |                          |
| 252.222.7002 | Compliance with Local Labor Laws (Overseas)                                                | Jan 1997       |                      |                          |
| 252.223.7008 | Prohibition of Hazardous Chemicals                                                         | Jan 2023       |                      |                          |
| 252.225.7056 | Prohibition Regarding Business Operations with the Xinjiang Uyghur Autonomous Region       | Jan 2023       |                      |                          |
| 252.225.7060 | Prohibition on Certain Procurements from the Xinjiang Uyghur Autonomous Region             | Jan 2023       |                      |                          |
| 252.232.7010 | Liens on Contract Payments                                                                 | Dec 2006       |                      |                          |
| 252.233.7001 | Choice of Law (Overseas)                                                                   | Jan 1997       |                      |                          |
| 252.237.7010 | Prohibition on Interrogation of Detainees by Contractor Personnel                          | Jan 2023       |                      |                          |
| 252.244.7000 | Subcontracts for Commercial Products or Commercial Services                                | Nov 2023       |                      |                          |
| 252.246.7004 | Safety of Facilities, Infrastructure, and Equipment for Military Operations                | Oct 2019       |                      |                          |
| 252.247.7023 | Transportation of Supplies by Sea                                                          | Oct 2024       |                      |                          |

FAR Clauses Incorporated by Full Text

52.212-4 Contract Terms and Conditions Required to Implement Statutes or Executive Orders-Commercial Products and Commercial Services, (Tailored) (Jan 2025)

Contract Terms and Conditions Required to Implement Statutes or Executive Orders-Commercial Products and Commercial Services (Jan 2025)

(a) The Contractor shall comply with the following Federal Acquisition Regulation (FAR) clauses, which are incorporated in this contract by reference, to implement provisions of law or Executive orders applicable to acquisitions of commercial products and commercial services:

(1) 52.203-19 Prohibition on Requiring Certain Internal Confidentiality Agreements or Statements (Jan 2017) (section 1403 of Division E, Title VII, of the Consolidated and Further Continuing Appropriations Act, 2013 (Pub. L. 113-234) and its successor provisions in subsequent appropriations acts (and as extended in continuing resolutions).

(2) 52.204-23 Prohibition on Contracting for Hardware, Software, and Services Developed or Provided by Rongevity Ltd Covered Entities (Dec 2023) (Section 1624 of Pub. L. 115-91)

Filters Sort

High Missing Clause [Dec 2025 Playback]

MISSING: 52.252-2 - Clauses Incorporated by Reference

Rationale

The contracting officer shall insert this clause in solicitations and contracts in order to incorporate clauses by reference.

Generated by Eudia

High Missing Clause [Dec 2025 Playback]

MISSING: 252.246-7003 - Notification of Potential Safety Issues

Rationale

Use this clause in solicitations and contracts, including solicitations and contracts using FAR part 12 procedures for the acquisition of commercial products and commercial services, for the acquisition of— (1) Repairable or consumable parts identified as critical safety items; (2) Systems and subsystems, assemblies, and subassemblies integral to a system; or (3) Repair, maintenance, logistics support, or overhaul services for systems and subsystems, assemblies, subassemblies, and parts integral to a system.

Generated by Eudia

High Incorrect Fill in Language [Filter Playback (NEW)]

Mark as Reviewed

**Figure 4 ACR: Augmented Contract Review**

Insights provides both document-level query capability and portfolio-level intelligence. The AskAI feature enables natural language queries against historical contract data, transforming institutional knowledge into an accessible resource. The dashboard provides real-





time visibility into portfolio health, obligation tracking, and risk identification. Together, these capabilities support both individual contracting actions and organizational decision-making.

The key structural principle across all three applications is that AI advises while humans decide. Every finding can be accepted, modified, or rejected based on professional judgment. Every recommendation includes explanation sufficient for the contracting officer to evaluate its validity. And every decision remains the responsibility of the warranted officer whose signature obligates the government.

**Research Question 3:** What measurable impact can this type of system produce when deployed in an operational environment?

Preliminary measurements indicate meaningful impact across multiple dimensions:

*Procurement Acquisition Lead Time Reduction.* Early indicators suggest that human-in-the-loop AI reduces several time-intensive components of the acquisition workflow, including compliance review, clause validation, and document preparation. Through testing, we estimate PALT reductions varying by acquisition type—with simpler actions such as Simplified Acquisition Procedures under \$350,000 showing potential reductions of [31%], while more complex Source Selection efforts show potential reductions of [43%]. Initial modeling of the Eudia Unified Platform deployed at scale projects a potential 34% reduction in average PALT across operational contracting squadrons. As additional base actions are processed through the system, we expect to further quantify reductions attributable to improved document accuracy, fewer rework cycles, and faster regulatory alignment.

*Improved Document Accuracy and Consistency.* Initial system logs show that the AI reviewer consistently detects missing clauses, misaligned citations, incomplete justification language, and regulatory conflicts in human-generated documents. Contracting personnel report that structured feedback and regulation-linked explanations help them correct issues earlier in the process when corrections are least costly. Early evidence indicates reduced error rates in key document types including market research reports, sole source justifications, and award documentation.

*Increased Compliance Confidence.* User feedback from 39 CONS indicates that contracting officers value clear, citation-based explanations outlining why a document section may create regulatory risk. This transparency improves confidence in both the final product and the decision-making process. Contracting officers report feeling more confident defending their work during peer review, legal review, and potential audit because they can trace every element to regulatory authority.

*Workforce Development Benefits.* Preliminary findings suggest the system supports junior contracting officers by reinforcing training concepts, surfacing relevant regulatory guidance, and demonstrating how experienced personnel reason through contract documentation. Rather than replacing professional development, human-centered AI accelerates it, providing real-time coaching that helps less experienced personnel develop judgment more quickly.

**Research Question 4:** How does the ingestion of historically awarded contract data influence contracting officer decision-making, compliance behavior, and workflow speed?

The Insights platform's ability to query historical contract data has demonstrated significant influence on contracting workflows. Contracting officers report that access to organizational precedent changes how they approach new actions. Rather than starting from scratch or relying on personal memory, they can quickly identify similar prior actions, review successful approaches, and adapt proven language to new circumstances.



This capability accelerates workflow speed by reducing research time from hours to seconds. It improves compliance behavior by making it easy to identify how similar situations were handled and whether those approaches withstood review. And it influences decision-making by providing contracting officers with broader context than any individual could retain, enabling more informed judgments about risk, approach, and documentation strategy. Equally important, portfolio-level visibility through the Insights dashboard enables leadership to make resource allocation decisions based on actual data rather than estimates. When leaders can see workload distribution, obligation rates, and risk concentration in real time, they can proactively address bottlenecks and prevent delays before they impact mission readiness.

## **Conclusion**

This research demonstrates that the fastest path to accelerating defense acquisition is not replacing human judgment but amplifying it. Human-centered AI designed to augment contracting officers rather than automate their work addresses the regulatory, audit, and accountability concerns that automation-first approaches create while delivering measurable improvements in speed, quality, and confidence.

The Eudia Unified Platform, validated through operational deployment at the 39th Contracting Squadron and informed by discovery across the acquisition enterprise, provides a model for how AI can be structured to support rather than supplant professional expertise. By keeping humans in control of every decision while providing them with capabilities that exceed what any individual could accomplish alone, this approach aligns with both the statutory structure of federal acquisition and the practical needs of an overworked workforce.

The implications extend beyond any single contracting squadron. The challenges identified in this research—regulatory complexity, disconnected systems, data-rich but insight-poor environments, and misaligned expectations between requirement owners and contracting officers—exist across the Department of War acquisition enterprise. Solutions proven at the squadron level have potential to scale, delivering PALT reductions that translate directly into faster delivery of warfighting capability.

## **Acknowledgments**

This research would not have been possible without the exceptional collaboration of acquisition professionals across the Department of War who generously contributed their time, expertise, and candid perspectives. We are deeply grateful to the men and women of the 39th Contracting Squadron at Incirlik Air Base, whose willingness to test operational prototypes and provide continuous feedback shaped every aspect of the Eudia Unified Platform. We extend particular thanks to Captain Cameron Cochran and Sergeant William Rife, whose partnership was instrumental to the success of this engagement.

We also acknowledge the invaluable contributions of stakeholders at Space Systems Command, the Office of the Chief of Space Force Contracting, the Air Force Test Center, and Headquarters Air Force Materiel Command. Their willingness to engage in guided discussions about workflow challenges, capability gaps, and opportunities for innovation provided essential context that strengthened this research.

We are fortunate to have had so many acquisition professionals at every level from flight to headquarters willing to share their experiences and perspectives for the purpose of advancing this work. Their commitment to improving defense acquisition, even amid demanding operational responsibilities, reflects the professionalism and dedication that characterizes the acquisition workforce. This research is ultimately in service of their mission: delivering warfighting capability to those who defend the nation.



THIS PAGE INTENTIONALLY LEFT BLANK



## PANEL 10. EMPOWERING THE PROFESSIONAL: SHAPING THE FUTURE OF ACQUISITION TALENT

Wednesday, May 6, 2026, 1445 – 1600 ET (1145 - 1300 PT)

### Panel Summary:

The success of any acquisition reform is fundamentally tied to the caliber and readiness of the professionals executing it. This executive panel brings together the Service Directors of Acquisition Career and Talent Management to discuss the strategic development of the Department's most critical asset: its people. By highlighting cross-service initiatives in education, credentialing, and certification, these leaders will explore how to synchronize efforts to build a modern workforce capable of accelerating warfighting capabilities in a complex global environment.

**Chair: Bob Marion, LTG USA, (ret.),** Professor of the Practice, Naval Postgraduate School; former Military Deputy, Assistant Secretary of the Army, Acquisition, Logistics and Technology (ASA ALT)

### Panelists:

**Army Acquisition Career Management** – *Ronald R. Richardson, Jr., Director, U.S. Army Acquisition Support Center and DACM*

**Navy Acquisition Talent Management** – *Renee King, Acting U.S. Navy Director, Acquisition Talent Management (DATM)*

**Air Force Acquisition Career Excellence** – *Ruben Rios, U.S. Air Force Director, Acquisition Career Management (DACM)*



**Bob Marion, LTG USA, (ret.)**—LTG Bob Marion began his career in the Acquisition Corps as an Assignment Officer in the Acquisition Management Branch, U.S. Army Personnel Command. He later served as the Assistant Project Manager (APM) for UH-60 Black Hawk A/L Production and Fielding, Utility Helicopter Project Office (UHPO), Aviation and Missile Command (AMCOM). He was then assigned to establish the Product Manager's (PdM's) Office for Black Hawk Modernization (UH-60M), UHPO, AMCOM, and served as the Acting PdM.

LTG Marion served as the Chief, Acquisition Management Branch in the Officer Personnel Management Directorate of the Army Human Resources Command and then deployed as the Assistant Secretary of the Army (Acquisition, Logistics and Technology) (ASA (ALT)) Forward Representative for the Deputy Assistant Secretary of the Army-Procurement. After changing command, LTG Marion served as the Assistant Deputy for Acquisition and Systems Management in ASA(ALT). He later served as the Program Executive Officer for Aviation until assigned as the Deputy for Acquisition and Systems Management, ASA(ALT). LTG Marion then deployed and served as the Deputy Commanding General for the Combined Security Transition Command - Afghanistan. He currently serves as the Principal Military Deputy to the ASA(ALT) and Director of the Army Acquisition Corps.

LTG Marion is a graduate of the Aviation Officer Basic and Advanced Courses, the Combined Arms Staff Services School, the Air Command and Staff College, and the Air War College. He also graduated from the Defense Acquisition University's Program Manager and Executive Program Management Courses. He earned a Master of Business Administration from George Mason University, a Master of Military Operational Art and Science and a Master of Science in Strategic Studies, both from the U.S. Air Force's Air University. He holds a Bachelor of Science in Labor Relations from the University of South Alabama.

His military awards and decorations include the Senior Aviator Badge, the Air Assault Badge, the Defense Superior Service Medal, the Legion of Merit, the Defense Meritorious Service Medal (1OLC), the Meritorious Service Medal (4OLC), the Army Commendation Medal, the Army Achievement Medal



(3OLC), the National Defense Service Medal, the Global War on Terrorism Expeditionary Medal, the Global War on Terrorism Service Medal, the Korea Defense Service Medal, the Army Service Ribbon, the Overseas Service Medal, the Afghanistan Campaign Medal, the NATO Medal, and the German Armed Forces Badge for Military Proficiency.



**Ronald R. Richardson, Jr.—** currently serves as the Director of the U.S. Army Acquisition Support Center (USAASC). In this role, he oversees the Army Acquisition Workforce (AAW) and supports the Army's Capability Program Executives in the areas of human resources, resource management, program structure, acquisition information management and program protection.

Mr. Richardson has over 30 years of medical, information and weapon system acquisition experience as both an Army civilian and a U.S. Army officer. Before coming to USAASC, he served as the Director of Acquisition and Operations for Program Executive Office (PEO) Soldier. Prior to joining PEO Soldier, he was the Deputy Project Manager for the DOD Healthcare Management System Modernization (DHMSM®) Program, a \$14B Major Automated Information System (MAIS) acquisition to replace the legacy Military Health System (MHS) Electronic Health Record (EHR) with an off-the-shelf (OS) system now known as MHS GENESIS. Before that, he was the Product Lead for Increment 3 of the Integrated Electronic Health Record (iEHR) Program in the DOD/Department of Veterans Affairs Interagency Program Office (IPO). Prior to joining the DOD/VA IPO, he served as the Director of Acquisition Review and Analysis for the Office of the Assistant Secretary of the Army, Acquisition, Logistics and Technology (ASA(ALT)). Before joining ASA(ALT), Mr. Richardson served in a multitude of Military, Civilian and Private Sector positions culminating in his selection for Senior Service College.

Mr. Richardson received his M.S. in Biomedical Engineering from Duke University, and his M.S. in National Resource Strategy from the Industrial College of the Armed Forces. He is also a graduate of the U.S. Army Command and General Staff College. He is the recipient of the Superior Civilian Service Medal (3), the Meritorious Civilian Service Medal (2), the Civilian Service Achievement Medal, the Army Staff Identification Badge and the Order of Military Medical Merit (O2M3). Mr. Richardson also holds multiple professional memberships and certifications, including Advanced Defense Acquisition Workforce Improvement Act (DAWIA) Certification in Program Management (previously Level III) and previous membership in the Army Acquisition Corps.



**Ms. Renee King—** is the Director, Acquisition Talent Management office under the Assistant Secretary of the Navy Research, Development and Acquisition (ASN RDA). Ms. King assumed this role in March 2024. In this role, she leads and synchronizes a staff of ten action officers and contractors, who support the Navy's 54,000-member Acquisition Workforce and the leadership and developmental programs and initiatives that enhance the capabilities of the Navy Warfighter.

In 2021, Ms. King competed and was selected to participate in the Naval Sea Systems Command (NAVSEA) Commander's Executive Fellows (CEFP) program as part of Cadre VII. This distinguished leadership initiative provided her with the chance to assist multiple commands, including the Program Executive Office (PEO) Ships, United States Fleet Forces Command, the NAVSEA Transformation Office (NTO), the Shipyard Infrastructure Optimization Program (SIOP) and Industrial Operations in various roles from an Acquisition Program Manager (APM) to Chief of Staff.

Prior to her selection into CEFP, Ms. Renee King served as the Director for the NAVSEA Acquisition Workforce Program Office in SEA 10. Ms. King led three enterprise-wide programs, the Defense Acquisition Workforce Development Account (DAWDA), the Defense Acquisition Workforce Improvement Act (DAWIA) Program and the Naval Acquisition Development Program (NADP). These programs collectively supported more than 20,000 Acquisition Workforce members, funded approximately \$18M in incentives and training to those employees (annually), and created a pipeline of more than 400 NADP Interns and Professionals to NAVSEA's acquisition workforce. In September 2016, Ms. King graduated





from the NAVSEA Journey Level Leader (JLL) Program, in which her 90- day rotation supported Surface Maintenance Engineering Planning Programs (SURFMEPP) Third Party Planning division, engaging in the technical management of three major contracts.

Ms. King particularly enjoyed her role at Trident Technical College, supporting veterans, as she is a United States Marine Corps veteran herself. She continues to support and volunteer her time with veterans and veterans-based programs. Ms. King is an active member of the Women Marines Association (WMA), Federally Employed Women (FEW), and Blacks In Government (BIG). Ms. King holds a Bachelor of Science degree in Business Management from Park University and is currently pursuing a master's degree in organizational leadership from Southern New Hampshire University.



**Mr. Ruben Rios**—is the Director of Acquisition Career Management Directorate, Assistant Secretary of the Air Force for Acquisition, Technology, and Logistics (SAF/AQH). Mr. Rios is responsible for the integrated management of the acquisition workforce across all functional areas. He provides acquisition human resources policy and strategic planning while managing the training and development of civilian and military acquisition personnel across the department's acquisition enterprise. Mr. Rios is also designated as the Career Field Manager for both military and civilian Scientists, Engineers, and Acquisition Program Managers. His team also provides personnel management services to the SAF/AQ Headquarters Staff and is the Department of the Air Force's Acquisition Demo Program Management Office.

Mr. Rios was an active-duty Air Force officer for 21 years, before retiring in 2010. He was commissioned through the Air Force ROTC program in 1988, after graduating with a bachelor's in mechanical engineering from the New Jersey Institute of Technology. His career in the DAF includes assignments within various System Program Offices, executing the roles of ground systems engineer for the Defense Satellite Program Mobile Ground Terminal System, flight test director for the E-767 Japanese AWACS, and Chief of Modeling & Simulation/Test for the E-10A/MP-RTIP program. His operational assignments include weapons controller/director, and later, senior director with the 963rd Airborne Warning and Control System (E-3 AWACS) Squadron. During his tenure with the 963rd, he participated in the closing phase of Operation Desert Storm, as well as numerous other deployments supporting operations Southern/Northern Watch, Provide Comfort, and Uphold Democracy.

His staff assignments include tours at HQ PACAF, where he served as a Pol-Mil officer and as an Action Officer in PACAF's Commander's Action Group. He served various tours at Air Staff, including, Branch Chief for Financial Execution/Budgeting within the Deputy Assistant Secretary of the Air Force for Acquisition Integration (SAF/AQX) and Deputy Division Chief for Program Integration with the Global Power Directorate (SAF/AQP). While at Air Staff, he was deployed to Islamabad Pakistan in support of Operation Enduring Freedom and was the Chief of Staff for the Office of the U.S Defense Rep. to Pakistan. Upon his return from Pakistan, he was selected to command the 364th Air Force Recruiting Squadron. In 2009, he returned for a second assignment with Air Staff, and assumed the position of Dep Chief, Program Integration Div., of the Global Power Directorate (SAF/AQP) before retiring.



THIS PAGE INTENTIONALLY LEFT BLANK



## PANEL 11. OPPORTUNITIES AND CONSIDERATIONS FOR ACQUISITION IN THE SECURITY COOPERATION SPACE

Wednesday, May 6, 2026, 1615 – 1730 ET (1315 - 1430 PT)

### Panel Summary:

As the Department of War (DOW) implements its new Acquisition Transformation Strategy, the mission space of security cooperation is undergoing a fundamental shift to enhance global deterrence. This panel explores the critical integration of allies and partners into the U.S. defense enterprise, addressing long-standing hurdles in Foreign Military Sales (FMS) and technology transfer. By examining new frameworks for consolidating global demand signals, leveraging allied R&D, and opening pathways for non-traditional and commercial technologies, these researchers provide a roadmap for building a more resilient, responsive, and collaborative international industrial base.

**Chair: Alan G. Gorowitz**, Senior Strategy Advisor, Defense Security Cooperation Agency (DSCA)

### Presenters:

**Improving the DoW's Ability to Leverage Allied and Partner Capabilities for Burden Sharing** – *Dr. Jennifer Moroney, RAND Corporation*

**Total Demand: Integrating Foreign Military Sales and Domestic Acquisition to Fortify the Defense Industrial Base and Strengthen Alliances** – *Aidan Winn and Jennifer Moroney, RAND Corporation*

**Embracing Non-Traditional Solutions: Evolving Security Cooperation for the Modern Threat Environment** – *Jarrett Lane, Institute for Defense Analyses (IDA)*



**Alan G. Gorowitz**—serves as Senior Strategy Advisor, providing analysis and recommendations regarding the DoD security cooperation enterprise in the context of DoD strategy, future scenarios, and emerging trends. He serves as a Senior Leader, supporting the Defense Security Cooperation Agency (DSCA) leadership team.

Mr. Gorowitz has served for over 20 years in the Department of Defense security cooperation community. He began his Department of Defense career in 2000 at the George C. Marshall European Center for Security Studies in Germany, supporting partner defense reform and strategy development through advising, conferences and seminars, and national security professional networks.

Returning to the U.S. in 2007, he served in the Office of the Secretary of

Defense (Policy) for over ten years providing policy advice and oversight including assignments as Director, Counterterrorism Partnerships; Chief of Staff and Director, Partnership Policy and Strategy; and Domestic Terrorism and Cyber Threat Information Sharing Coordinator. In 2017, he was detailed from Policy to DSCA to stand up a Planning and Program Design team.

He joined DSCA in 2018 as Assistant Director for Strategy, Plans, and Policy (SPP), facilitating strategic planning, leading security cooperation administration and execution policy, and promoting comprehensive security cooperation planning and program design.

Prior to Federal Service, Mr. Gorowitz served at the Center for International Strategy, Technology, and Policy (CISTP), Sam Nunn School of International Affairs, Georgia Institute of Technology.

Mr. Gorowitz is a graduate of the National War College. His awards include the Office of the Secretary of Defense Award for Exceptional Service (3), the OSD Achievement Award, and the Department of the Army Award for Superior Civilian Service.



# **Improving the DoW's Ability to Leverage Allied and Partner Capabilities for Burden Sharing**

**Dr. Jennifer DP Moroney**—The RAND Corporation

**Karen Schwindt**—The RAND Corporation

## **Abstract**

The transition of DSCA and DTSA to OSW A&S is intended to address long-standing concerns that the FMS/DCS system is bloated and sluggish, often delivering defense products to partners and allies far too slowly. The transition is expected to support FMS and DCS reform by linking these processes, systems, and personnel more directly to DOW acquisition strategy. Keeping in mind that Security Cooperation is more than defense sales introduces the prospect of reaping additional benefits from this transition, pertaining to increasing burden sharing with key partners and allies and building more resilient supply chains that include allied and partner investments. This paper will identify and unpack opportunities that might otherwise be overlooked. For example, bringing together FMS/DCS Title 22 activities and international armaments cooperation under Title 10 (to include co-design/development/production/sustainment) could generate cases (that is, product transfer to allies and partners) that share aspects of both sales and R&D. The new FMS/DCS system could be designed to better enable this kind of hybrid case as well as other new ways to simultaneously meet US and allied needs in support of the warfighter. There are already a few examples of such hybrid cases, which this paper will analyze to generate potential lessons. Taking a step further, the new FMS/DCS system could also identify avenues to import innovative capabilities from allies and partners in a way that directly supports the DIBs of both the US and its allies and partners. The paper will integrate insights from a recent RAND publication on Exploring the "Optimal Pathway" to AUKUS Pillar II Success, which has specific lessons and best practices that can inform these changes.

## **Objectives and Purpose**

The paper will present and begin to unpack opportunities for working in new and innovative ways with key allies and partners that are linked to the recent changes per the new Acquisition Transformation Strategy (ATS; Office of the Under Secretary of Defense for Acquisition and Sustainment, 2025), including ways to build resilience in U.S. supply chains (Executive Order No. 14017, 2021) and opportunities for importing innovative ideas, technologies, and capabilities from key allies. It will also point to ways to operationalize burden sharing that drive Foreign Military Sales (FMS) and cooperative development opportunities that enable allies and partners to contribute directly to broader resilience and deterrence national security objectives.

The objective of this short paper and presentation is to, first, present a collection of tangible options to further burden sharing with allies, drawing upon completed RAND research and ideas flowing from the results of those research projects that could fit within the new ATS. The second objective is to solicit feedback from acquisition and security cooperation experts on these ideas to highlight the best options, which could be further studied.

The paper begins with a statement of the problem and background relevant to the problem to set the context. The background draws on RAND's research that relates to trends we are seeing in the way allies and partners desire to work with the U.S. and how this is relevant to the new ATS, as implementation is evolving. Then, the paper briefly identifies six options that are meant to help inform the implementation of the ATS. Each option is presented with a short set of researchable questions.



## Problem Statement

Persistent misalignment between acquisition and security cooperation systems, driven by institutional, statutory, and procedural barriers, has limited the United States' ability to deliver capabilities to allies and partners at the speed and scale required to sustain deterrence and resilience in an era of strategic competition. While the United States has long collaborated with allies and partners on capability development, these efforts have not been systematically integrated into acquisition strategy or optimized for speed, scalability, and industrial resilience (Executive Order No. 14017, 2021). Despite repeated reform initiatives, coordination between acquisition and security cooperation systems has been hindered by structural, statutory, and cultural barriers. Separate legal authorities under Title 10 and Title 22 divide responsibility between the Department of War (DoW) and Department of State (DoS), creating parallel processes with distinct oversight, funding streams, and priorities (Sargent, 2023). DoW acquisition programs are planned through multiyear budget cycles focused on U.S. force requirements, while FMS cases are reactive to partner demand and subject to DoS review (GAO, 2023a). These misaligned timelines and authorities have made it difficult to synchronize production, co-development, or sustainment with allies, undermining the ability to build collective capacity and maintain deterrence.

Institutionally, acquisition and security cooperation communities operate with different professional cultures and information systems. Acquisition personnel are trained in contracting and program management, whereas security cooperation professionals focus on diplomacy, partner engagement, and export control (RAND Corporation, 2019). Limited cross-training and incompatible data systems have reinforced stovepipes, producing redundant reviews and inconsistent demand forecasting (GAO, 2020). As a result, even when the DoW and DoS share strategic objectives, they lack the mechanisms to jointly plan, fund, and execute programs that incorporate allied participation from the outset. These systemic barriers, not lack of intent, explain why past integration efforts have yielded limited results (GAO, 2021).

U.S. planning and acquisition processes have often defaulted to a model in which the United States leads innovation, and partners primarily consume capabilities through FMS or other transfer mechanisms. While many U.S. programs have followed a supplier–consumer model, notable exceptions, including the F-35 Joint Strike Fighter program, demonstrate that co-development and deeper integration is possible but remain the exception rather than the norm (Gertler, 2020). Although this approach has delivered interoperable capabilities, it is increasingly misaligned with the requirements of operating in a competitive strategic environment, where deterrence, resilience, and burden sharing are essential to sustaining advantage (Defense Security Cooperation Agency, 2023; DoW, 2026a).

Alternative ways of thinking about allied contributions to U.S. and coalition warfighter capabilities are essential. Allies and partners are no longer solely capability recipients—they are increasingly sources of innovation, niche industrial capacity, and operational expertise (Saum-Manning et al., 2022). Many are investing in sovereign capabilities, diversifying supply chains, and seeking reciprocal relationships that include co-development and co-sustainment. Simultaneously, global demand for defense systems, particularly space capabilities and advanced technologies, continues to outpace U.S. production capacity, highlighting structural limitations in the current acquisition and security cooperation enterprise that constrain the nation's ability to sustain deterrence and technological advantage (DoD, 2023; GAO, 2023b).

Recent reforms, including the new ATS and the realignment of DSCA and DTSA under the Office of the Under Secretary of War for Acquisition and Sustainment (OUSW[A&S]), present important opportunities to address these challenges (DoW, 2026b). These changes signal a recognition that security cooperation, technology security, and acquisition must be more tightly integrated to meet emerging operational demands. Bringing FMS and direct commercial





sales (DCS) processes into closer alignment with DoW acquisition and allied R&D frameworks is an important step, but how this integration is implemented—particularly in terms of how authorities are delegated, how interagency coordination is managed, and whether new processes avoid duplicating existing ones—will determine whether the intent behind the ATS is realized. Additionally, any integration effort must preserve the essential checks and balances that Titles 10 and 22 were designed to maintain, ensuring that efficiency gains do not compromise policy oversight or technology security.

## **Policy Alignment**

These challenges and opportunities emerge within a policy environment that strongly supports acquisition transformation, allied burden sharing, and industrial resilience. Recent strategies, executive orders, and congressional actions collectively underscore the need to align acquisition, security cooperation, and technology–security functions to sustain deterrence and resilience in an era of strategic competition.

Recent executive orders and the Presidential Memorandum on the National Defense Industrial Strategy (DoD, 2024) direct and operationalize the current policy foundation for defense industrial and supply chain resilience. Executive Order No. 14017 (2021) specifically directed a whole of government effort to expand domestic and allied production capacity while Executive Order No. 13806 (2017) created the Defense Industrial Base Council and modernized technology release processes.

At the congressional level, the House Foreign Affairs Committee (HFAC) has held hearings and issued statements emphasizing the need for faster, more transparent FMS processes and greater allied burden sharing, echoing executive branch priorities and providing legislative momentum for reform. Additionally, recent discussions in congressional hearings have highlighted the potential elevation of the Deputy Assistant Secretary of War for International Cooperation to an Assistant Secretary–level position to strengthen allied co-development leadership (House Foreign Affairs Committee, 2023).

Within the DoW, the ATS and the 2024 Defense Industrial Base Strategy directly advance priorities outlined in the 2026 National Defense Strategy (NDS), which defines deterrence, resilience, and allied burden sharing as core U.S. advantages. Together, these policies provide the strategic and institutional foundation for the options and implementation pathways examined in this paper.

## **Background and Context**

The transition of DSCA and DTSA under the OUSW(A&S) marks a significant structural reform intended to address long-standing concerns that the FMS system is overly complex and slow. By aligning security cooperation and technology security functions more closely with acquisition authorities, this reform creates an opportunity to synchronize requirements, acquisition processes, and partner demand signals. The realignment, announced in December 2022, reflects a recognition that acquisition, technology security, and security cooperation must operate as mutually reinforcing elements of a single enterprise to sustain deterrence and resilience.

Recognizing that Security Cooperation encompasses more than defense sales introduces the prospect of additional benefits from this transition, such as increasing burden sharing with key partners and allies, building more resilient supply chains that include allied and partner investments, and leveraging partner R&D in niche areas to address known capability gaps and expand industrial capacity. If effectively implemented, these changes could enable a



shift from a transactional model of security cooperation toward a more integrated approach that supports shared capability development, industrial resilience, and collective deterrence.

This short paper identifies and unpacks several opportunities, some of which might otherwise be overlooked, to operationalize these reforms. It integrates insights from recent RAND publications focused on FMS reform (Saum-Manning, 2025), allied and partner sovereign capabilities, third country suppliers, and the AUKUS partnership (U.S. President et al., 2021), which all offer lessons and best practices that can inform implementation of the ATS and related initiatives.

## Responding to Allied and Partner High-Level Trends

A RAND research project conducted in 2023, *Insights on U.S. Ally and Partner Views of Strategic Competition: Implications for the Department of the Air Force*, explored how allies and partners prefer to engage with the United States in security cooperation. The study identified two notable shifts from prior decades: first, a growing emphasis on developing sovereign capabilities and, in some cases, a preference for *minilateral collaboration*, particularly around technology cooperation, to achieve mutual objectives more efficiently.

Although the pursuit of *sovereign capabilities* is not new, many countries have become more explicit in articulating how they intend to collaborate with foreign partners. This includes revising national defense strategies and procurement policies to balance domestic production, acquire U.S. systems, and purchase from other suppliers. For this analysis, sovereign capability is defined as *the ability to control the development, production, sustainment, or employment of a nation's defense assets*. RAND's research found that both pragmatic and ideological factors drive this pursuit—some nations seek to meet unique operational requirements or improve supply chain reliability, while others aim to enhance affordability and generate economic benefit (Executive Order No. 14017, 2021). Defense industrial base (DIB) considerations often factor heavily into a country's sovereign capability motivations.

Political pronouncements, strategic reviews, and defense white papers reveal that the eight countries examined in this study generally follow one or more of three approaches:

1. Develop broad sovereignty, seeking to become fully or largely self-sufficient (Brazil, France, India, Turkey)
2. Maintain sovereign capability in specialized or core areas while engaging internationally in other areas (Australia, France, South Africa, Sweden, Thailand)
3. Use diverse suppliers for defense assets to reduce dependence on any single country or system and preserve strategic autonomy (Thailand, Turkey).

We continue to see this trend evolve. In February 2026, Canadian Prime Minister Mark Carney announced Canada's new "Buy Canadian" defense strategy, which aims to reduce Canada's dependency on the United States while revitalizing an otherwise most dormant DIB (Government of Canada, 2026). He plans to add 25,000 new jobs in the defense sector in the next decade. The long-awaited strategy commits to deepen partnerships with Europe and key Indo-Pacific allies in response to Trump's aggressive security and trade posture. He was careful to defend the NORAD relationship with the United States (Blanchfield, 2024).

The growing emphasis on sovereign capability provides essential context for the options presented in the following section. As the DoW implements the new ATS, it should aim not only to streamline processes but also to ensure that acquisition and security cooperation frameworks effectively incorporate allied capabilities, align with partner incentives, and support shared strategic objectives.



The convergence of these institutional reforms and allied trends underscores both the urgency and opportunity for the DoW to rethink how it engages partners across the acquisition and security cooperation enterprise. The following section builds on this momentum by identifying specific options for creating more flexible pathways that link FMS, DCS, and international armaments cooperation while protecting America's technical edge. Each option is designed to translate the policy intent of the ATS into practical mechanisms that enhance allied participation, strengthen industrial resilience, and sustain collective deterrence.

### **Creating Flexible Pathways: Options for Consideration Moving Forward**

This section identifies six options for DoW consideration. While these options were developed using completed RAND research, these options have not been analytically evaluated for sufficiency or effectiveness. As such, each option concludes by identifying areas for further research.

#### **Creating Options for Working with Allies: Mixing FMS and International Armaments Cooperation Deliberately**

With the realignment of DSCA under OUSW(A&S), the two authorities and funding streams of Title 22 (FMS) and Title 10 (International Armaments Cooperation) now fall under a single oversight structure. This alignment enables the DoW to better coordinate funding timelines, contracting mechanisms, and technology release decisions that were previously managed through separate channels. By integrating these processes, the DoW can design hybrid pathways that allow partners to participate in early stage R&D, contribute to co-production or co-sustainment, and procure end items within a single program format. Such models would expand the industrial base, thereby supporting U.S. allies and forces, while also advancing NDS 2026 objectives of deterrence, resilience, and allied burden sharing. This hybrid approach moves away from the historic approach that treated these mechanisms as mutually exclusive.

A useful reference point is the F-35 Lightning II Joint Strike Fighter program, which blends cooperative development among core partners with FMS cases for additional participants. While not structured as a hybrid case, the program demonstrates how varying levels of partner participation, ranging from co-development and industrial contribution to traditional procurement, can coexist within a single capability ecosystem. The F-35 program also demonstrates how allied participant serves as a way to scale the cost of developing and producing these aircraft, thereby creating an affordable way for partners to contribute to or acquire an advanced capability while also bolstering the U.S. industrial base (Congressional Research Service, 2026).

Building on this precedent, the Department could explore more intentional hybrid case constructs in which allies and partners are offered structured options to

- participate in early stage R&D or prototyping efforts (Title 10)
- contribute to co-production or component manufacturing
- procure end items or additional quantities through FMS (Title 22)

This approach would better align with partner preferences for sovereign capability development while also expanding the industrial base supporting US and allied force requirements. Hybrid approaches could also improve demand signaling, increased production scale, shared development costs, and more resilient supply chains. However, challenges relating to legal authorities, funding streams, program management, and oversight, particularly with regard to congressional notification and DoS review, raise questions about how hybrid constructs would be structured and governed in practice.



### ***Key Questions for Further Research***

- What are the risks for encouraging hybrid cases? Are there any efficiencies to be gained, or would this muddy the authorities' pool in a way that is unhelpful?
- What new administrative processes or arrangements might have to happen to enable this?
- What oversight mechanisms at the DoS and DoW are needed?

### **Multinational FMS for Europe and Beyond**

The 2026 NDS emphasizes the importance of reviving U.S. industry and mandates that allies assume primary responsibility for conventional defense in their respective regions. However, the U.S. security cooperation enterprise remains optimized for bilateral transactions. The U.S. DIB currently receives low-volume signals from dozens of individual partners rather than consolidated demand signals. Without a mechanism to aggregate requirements, manufacturers cannot achieve the economies of scale necessary to “supercharge” production capacity as directed by NDS Line of Effort 4.

To be sure, NATO is already doing some multinational FMS on a small scale to try and meet its most recent set of capability targets. However, many NATO allies face significant shortfalls in their purchasing power; and while almost all allies have committed to increased defense spending, their individual budgets remain insufficient to procure high-end, critical capabilities (e.g., integrated air and missile defense, long-range fires) at the scale required (NATO, 2023). Furthermore, the stovepiped nature of procurement efforts—rather than multinational, multiyear procurements—leads to a heterogeneous force structure that is difficult to sustain and lacks true interoperability. Despite initiatives like NATO’s Defense Production Action Plan (DPAP), allies are still pursuing unilateral solutions more frequently than multinational procurements (NATO, 2026). This fragmentation risks leaving critical requirements—such as the significant interest in counter-UAS and loitering munitions identified in the 2025 REPEAD cycle—unaddressed. The U.S. FMS system is often perceived as a barrier to, rather than an enabler of, such multinational collaboration.

In the Indo–Pacific theater, the absence of a formal alliance procurement infrastructure, comparable to NATO’s Support and Procurement Agency, continues to reinforce reliance on bilateral U.S. support. This dynamic places the United States in a position where it must often assume the principal burden of equipping partners during crises. Managing hundreds of discrete FMS cases without clear prioritization or administrative streamlining consumes significant bandwidth within OUSW(A&S), diverting attention from high-priority modernization efforts in both the European and Indo–Pacific regions. Persistent supply constraints and extended delivery timelines for key systems—such as Patriot air-defense units and loitering munitions—further erode deterrence and sustain dependence on U.S. stockpiles. Additional research is warranted to identify mechanisms that could improve regional coordination, accelerate delivery, and strengthen collective resilience.

### ***Key Questions for Further Research***

- Are there opportunities for economies of scale for multinational FMS that would directly benefit the United States? To what extent could a multinational FMS model measurably reduce administrative lead times and unit costs particularly for high-demand munitions and non-program of record (NPOR) systems?
- Could multinational FMS serve as a conduit to co-development and co-production opportunities to further enhance burden sharing opportunities?



- How can multinational FMS facilitate burden-shifting in the Indo–Pacific where formal alliance structures like NATO are absent?
- What are the primary barriers (e.g., regulatory and timeline issues) preventing the widespread adoption of multinational FMS for co-production and co-development efforts?

### **Expanding Trusted Partnerships Models to Key Allies**

For the purposes of this analysis, trusted partnerships refer to structured, institutionalized arrangements that enable select allies to participate across multiple phases of the acquisition lifecycle. AUKUS represents a significant test case in operationalizing high-trust defense cooperation. The partnership between Australia, the United Kingdom, and the United States demonstrates that when partners share aligned strategic objectives and robust technology protection frameworks, barriers to co-development of the most sensitive technologies (e.g., nuclear-powered submarines, quantum technologies, undersea warfare capabilities, and artificial intelligence) can be reduced through tailored governance and shared risk mechanisms.

However, the initiative also underscores the complexity of sustaining such an arrangement as it requires dedicated resources, clear decision-making structures, and consistent industry engagement (Moroney, 2022). Despite these ongoing issues, AUKUS has become the beacon for technological minilateralism, helping drive meaningful reforms to ITAR, and can provide many lessons for what to do and what not to do when it comes to trusted partnerships among the United States and its key allies. Expanding this model to other domains, such as space, would advance the NDS 2026 objectives of deterrence and resilience by accelerating capability integration and diversifying allied industrial participation. The DoW and DAF could examine how domestic-preference policies, export controls, and security cooperation objectives intersect in the space sector to inform future decisions on co-production/sustainment of nonsensitive components. Incremental steps for vetted partners to contribute could expand production throughput and stabilize supply chains while maintaining technology protection.

### **Key Questions for Further Research**

- What criteria are helpful to determine whether an ally should be in a “trusted partnership” with the United States? Preliminary criteria could include alignment of strategic objectives, ability to protect sensitive technologies, etc.
- What lessons and best practices can be learned from AUKUS pillar 2 that can be used to inform the creation of new minilateral trusted partnerships?
- Should the DoW look only at capable space allies, or wider, assuming a low level of sensitivity with regards to the hardware and software in question?
- Should the DoW consider a “pilot” trusted partnership construct in space, or other high priority capability areas to assess feasibility and scalability?

### **Alternative Options to FMS: How and When Could DCS Provide Greater Efficiencies?**

As the DoW modernizes security cooperation, a key question is not whether to replace FMS but how to optimize the mix of FMS and DCS (where allies and partners go direct to U.S. companies for their defense procurements rather than via the DoW) pathways based on capability type, partner needs, and timelines. Because DCS allows direct contracting between foreign customers and U.S. firms, it can shorten administrative timelines and leverage commercial production lines already operating at scale. In the space sector, where allied demand has surged and commercial innovation is robust, DCS could possibly provide faster access to dual-use technologies. However, the absence of a centralized DCS reporting





mechanism limits the DoW's visibility into aggregate export activity, complicating technology security oversight and industrial base planning.

The U.S. Space Force (USSF) is experiencing rising allied and partner demand for space capabilities. Between 2023 and 2024, the USSF's Space Systems Command (SSC) reported a 500% increase in requests for space capabilities through the FMS process, expanding beyond the capacity of the USSF to meet the demand (Albon, 2024). In parallel, the expanding commercial sector offers a host of dual-use technologies from a new generation of companies that can help satisfy allied and partner demand. It is not clear that FMS is the right approach for all capabilities or for all partners. This surge has overwhelmed the DAF and USSF in the near term, with case management, export control and ITAR checks all requiring excessive manpower and expertise.

The space sector offers a promising environment to examine greater reliance on the DCS model, yet several gaps warrant further study. Key uncertainties include how DCS could deliver measurable gains in speed and commercial integration without compromising technology–security oversight, and how the absence of centralized DCS tracking affects visibility into total U.S. space-cooperation activity, industrial-base throughput, and technology-transfer patterns. Addressing these questions through targeted research would clarify where and how DCS can complement FMS to meet allied demand more effectively.

### ***Key Questions for Further Research***

- Should DCS be formally endorsed by the DoW, and in what sectors and instances might this make sense? What are the risks and opportunities?
- When does DCS provide more or faster capability to allies and partners than FMS in space? Under what circumstances?
- How are the U.S. DIB-primes and commercially driven small-to-medium enterprises responding to the increase in space FMS and DCS cases? Which mechanism do they prefer and why? What are the costs, risks, and timeline trade-offs from the industry perspective?
- To what extent could the DAF and the U.S. DIB benefit from changes to the way DCS is managed and tracked? Could industry absorb some of the administrative costs for DCS space cases?
- What are the technology protection, sustainment, training, and interoperability implications of shifting from FMS to DCS for key space capabilities?

### **Leveraging Allied Capabilities and Innovations for Burden Sharing**

This option focuses on reverse integration—systematically incorporating allied and partner-developed capabilities into U.S. acquisition programs and force design. Building on previous options, the new FMS system could also identify avenues to import innovative capabilities from allies and partners in a way that directly supports the DIBs of both the United States and its allies and partners. This could involve adopting allied-developed systems, integrating subsystems into U.S. platforms, or leveraging partner-provided data and services.

While the DoW historically has much greater resources to put towards innovation, allies and partners often innovate more quickly and more efficiently out of necessity. In an environment where capability development timelines are compressing, leveraging allied innovation can provide faster pathways for fielding relevant capabilities. Examples of such innovation include autonomous systems, electronic warfare tools, and small satellite or space payload technologies developed by close allies. In the 2020s, U.S. allies and partners also have developed more commercial and scientific capacity than they did during the Cold War



suggesting new opportunities. In short, the DoW should be on the lookout for these niche capabilities and actively engaging key allies to see what they have and are working on.

Moreover, the acquisition and SC communities should be discussing specific ways to include key allies into the earlier stages of DoW acquisition and capability development processes. Previous RAND research has pointed to the need to build in exportability by design principles to DoW planning, so that allied and partner interests are captured and export control considerations are included from the start. It is not helpful to anyone to have allied interests strictly as an afterthought. It could be useful to institutionalize exportability in space acquisitions (Federal Register, 2025). The DoW should consider ways to embed exportability by design principles across acquisition lifecycles, from requirements definitions through contracting phases. Reviewing how ITAR, disclosure, and release policies interact with program design and systems could highlight opportunities for earlier partner integration, tested through mechanisms like the International JROC or other exportability working groups.

### ***Key Questions for Further Research***

- Could a framework be designed for understanding when and where the DoW should seek to leverage allied innovations? What criteria would be appropriate?
- What sectors are the highest priority?
- Are there examples for the DoW importing innovative capabilities from allies and partners that help to inform the path going forward?
- Are there any successful examples of exportability by design principles applied with successful outcomes?

### **Developing a Cross-Functional Workforce to Include Acquisition, International, and Technology-Release Functions**

The DoW workforce currently lacks the cross-functional integration needed to effectively leverage allied contributions, manage hybrid FMS/DCS models, or embed exportability by design principles across acquisition programs. Distinct career tracks and training pipelines have historically limited collaboration among acquisition, international affairs, and technology security professionals, resulting in fragmented decision-making. Development of a cross-functional workforce, combining expertise in acquisition, international engagement, export control, and system integration, may accelerate capability delivery and improve coordination of technology release decisions. For space specifically, integrating acquisition, international engagement, and technology-release functions could streamline FMS and DCS case management, enable faster partner integration, and enhance protection of sensitive technologies. Training programs could also include rotational assignments across SAF/IA, SAF/SQ, SSC, DSCA, DTSA, etc., to foster shared expertise and cross-organizational awareness. Strengthening collaboration among SAF/IA, SAF/SQ, SSC, DSCA, and DTSA and updating curricula through WAU and DSCU could help generate the expertise required to transform space FMS into a driver of resilience, deterrence, and strategic advantage to enable a globally integrated space enterprise (Space Force, 2025). This initiative aligns with both Executive Order No. 14017, Executive Order No. 13806, and the 2024 National Defense Industrial Strategy (NDIS), which emphasize workforce readiness and cross-section collaboration as pillars of industrial resilience. Performance metrics such as reduced case-processing times and increased partner participation could help evaluate its effectiveness.



### **Key Questions for Further Research**

- Who should take the lead for training: WAU for acquisition or DSCU for security cooperation? Or should an inter-school effort be created to combine knowledge and expertise?
- How large should this cadre be, and how should it be expected to grow over time?
- Who should take the lead in managing this cross-functional workforce: DSCU, joint office, or other?
- How large should this cadre be initially, and what is the projected growth trajectory over 5 to 10 years?
- What training programs and rotational assignments are needed to build the required expertise?
- How should performance and outcomes be measured to ensure that the workforce improves FMS/DCS integration and allied cooperation?

### **Allies and Partners Have Choices Too**

Recent RAND research entitled *U.S. Options for Identifying Third-Party Suppliers to Meet Ally and Partner Capability Needs* explored the conditions shaping the international defense market. We identified cases where U.S. allies and partners are looking to non-U.S. suppliers to meet their defense needs, and we wanted to understand, from the allied and partner perspective, why this is the case. The research found that the United States may not always be able to provide timely or tailored solutions for partners. In some cases, U.S. systems or processes are not optimally suited to address specific operational requirements. It was also clear in some instances that the U.S. system in question was not ideally suited to address a partner's specific needs. Third-party suppliers could be incorporated into a U.S. strategy as complementary options, ensuring that partners acquire needed capabilities more efficiently while still aligning with U.S. strategic objectives. The work highlights the variety of factors that could make a U.S. solution untenable, including such considerations as cost, timeline for availability, the functional features for the intended operational context, export control policies, and strategic competition.

The bottom line is that U.S. solutions may not always be optimal for partner needs. The United States can take steps to be more aware of these third-party suppliers and even support them, when it makes sense, because allies and partners do have a choice.

At the same time, allied industries can be encouraged—and in some cases incentivized—to establish a presence in the United States, particularly when they possess niche capabilities that complement U.S. defense-industrial needs. Such investment would expand domestic production capacity, diversify supply chains, and strengthen collective resilience. Recent policy initiatives, including Executive Order No. 14110, Executive Order No. 14114, and the NDIS 2024, explicitly call for allied co-production and industrial-base integration as pillars of deterrence and competitiveness. Secure foreign investment frameworks such as CFIUS and programs under the defense production act Title III provide mechanisms to facilitate these partnerships while safeguarding technology and national-security interests. Clarifying demand signals and establishing transparent pathways for allied participation would help translate these policies into tangible industrial-base outcomes.

### **Conclusion**

This paper has outlined actionable options for engaging allies and partners in new and innovative ways consistent with the objectives of the ATS. Each option warrants deeper analysis



to assess its feasibility, benefits, and potential risks. By identifying and operationalizing burden sharing approaches that integrate FMS, cooperative development, and allied innovation, the DoW can strengthen collective resilience, enhance deterrence, and advance shared national security objectives. Sustained implementation of these reforms, supported by clear policy guidance, coordinated governance, and continued engagement with industry and allies, will be essential to realizing the full potential of acquisition transformation. In short, we all can win.

## References

- Albon, C. (2024, October 28). *U.S. Space Force sees surge in foreign military sales demand*. Defense News.
- Baker, Sinéad. (2026, March 29). *NATO officials say members aren't buying weapons together, and it's a mistake*. Business Insider.
- Blanchfield, M. (2024, February 17). *Canada defense plans aim to reduce dependency on U.S.* Politico.
- Congressional Research Service. (2020, April 29). *F-35 Joint Strike Fighter (JSF) program* (CRS Report No. RL30563).
- Defense Security Cooperation Agency. (2023). *DSCA strategic plan, 2023–2027*.
- DoD. (2023, June 13). *Department of Defense unveils comprehensive recommendations to strengthen Foreign Military Sales* [Press release].
- DoD. (2024, January). *National Defense Industrial Strategy*.
- DoD Office of Inspector General. (2020). *Evaluation of the DoD's management of the Foreign Military Sales process* (DoDIG-2020-084).
- DoW. (2026a). *2026 National Defense Strategy*.
- DoW. (2026b). *Department of War finalizes realignment of the Defense Security Cooperation Agency and the Defense Technology Security Administration* [Press release].
- Executive Order No. 13806, 82 Fed. Reg. 34597 (2017, July 21). *Assessing and strengthening the manufacturing and defense industrial base and supply chain resiliency of the United States*.
- Executive Order No. 14017, 86 Fed. Reg. 11849 (2021, February 24). *America's supply chains*.
- Federal Register. (2025, April 15). *Reforming defense sales: Reforming foreign defense sales to improve speed and accountability*.
- GAO. (2021). *Defense acquisition reform: DoD efforts to streamline processes and integrate partners* (GAO-21-329).
- GAO. (2023a). *Foreign Military Sales: DoD should take additional actions to improve timeliness of case processing* (GAO-23-106188).
- GAO. (2023b). *GPS modernization: Space Force should reassess requirements for satellites and handheld devices* (GAO-23-106018).
- Gertler, J. (2020, April 29). *F-35 Joint Strike Fighter (JSF) program* (CRS Report No. RL30563). Congressional Research Service.
- Government of Canada. (2026, February 17). *Prime Minister Carney launches Canada's first Defence Industrial Strategy to strengthen security, create prosperity, and reinforce strategic autonomy*.
- House of Representatives, Committee on Foreign Affairs, Subcommittee on Oversight and Accountability. (2023, September 19). *Hearing on reforming the foreign military sales process*.
- Moroney, J. D. P. (2022, March). *Making AUKUS work*. The Strategist.



- Moroney, J. D. P., et al. (2024, September 13). *Insights on U.S. ally and partner views of strategic competition: Implications for the Department of the Air Force*. RAND Corporation.
- NATO. (2023, July 11–12). *Defence production action plan*. 2023 NATO Vilnius Summit, Vilnius, Lithuania.
- Office of the Under Secretary of Defense for Acquisition and Sustainment. (2025, November 7). *Acquisition transformation strategy: Rebuilding the arsenal of freedom*. DoD.
- RAND Corporation. (2019). *Reforming security cooperation: Integrating Title 10 and Title 22 authorities* (Report No. RR-2828).
- Saum-Manning, L. (2025). *Waiting to win: The FMS accountability gap*. RAND Corporation.
- Saum-Manning, L., Marquis, J. P., Chindea, I. A., Elinoff, D., Pigott, T., & Brennan, E. (2022). *Improving the timeliness of FMS* (Report No. RRA-1340-1). RAND Corporation.
- Sargent, J. F., Jr. (2023). *Foreign military sales: Process and policy* (CRS Report No. R4634). Congressional Research Service.
- Space Force. (2025, June). *U.S. Space Force international partnership strategy: Strength through partnership*.
- U.S. President, Prime Minister of Australia, & Prime Minister of the United Kingdom. (2021, September 15). *Joint leaders' statement on AUKUS*.





# **Total Demand: Integrating Foreign Military Sales and Domestic Acquisition to Fortify the Defense Industrial Base and Strengthen Alliances**

**Aidan Winn**—is a policy researcher at RAND specializing in security cooperation, European security, alliances, and U.S. arms transfers policy. Prior to RAND, she supported the Office of the Under Secretary of Defense for Policy, managing EUCOM security cooperation and contributing to National Defense Strategy implementation, partner capability development, and civilian harm mitigation initiatives. Previously she worked at Booz Allen Hamilton supporting the U.S. Air Force on strategy and partner engagement. Winn was a U.S. Army Center for Military History dissertation fellow, a Truman National Security Fellow, and a security studies instructor at UMass Lowell. She holds a PhD in war studies from King's College London.

## **Abstract**

This paper examines the historical structural separation between U.S. domestic defense acquisition and foreign military sales (FMS) and its impact on the defense industrial base. Traditionally, domestic procurement has been planned through internal Department of War processes, while allied demand is generated through partner-initiated FMS requests, resulting in fragmented and reactive demand signals for industry. As allied procurement has become an increasingly significant share of total demand, this bifurcated system has contributed to production instability, supply chain fragility, and delayed capability delivery.

To address these challenges, the paper proposes a “Total Demand” framework that integrates domestic and allied procurement into a unified planning construct. Leveraging advances in artificial intelligence and data analytics, this approach aggregates data from program planning, partner requests, readiness indicators, and industrial capacity to generate predictive demand forecasts. The analysis demonstrates how integrated demand signals can stabilize production, enable economies of scale, and strengthen supply chain resilience.

The paper concludes that adopting a Total Demand framework would enhance industrial readiness, improve the efficiency of defense exports, and reinforce alliance interoperability, positioning the United States to better support coalition operations in an era of strategic competition.

## **Introduction: The Bifurcated Demand Dilemma**

For decades, the United States has managed domestic military acquisition and foreign military sales (FMS) through two largely separate institutional systems. Domestic procurement is planned through the Department of Defense’s internal resource allocation mechanisms, while allied demand emerges through partner-initiated FMS requests. As a result, foreign partner demand signals frequently emerge outside the planning horizons that guide domestic production decisions. These parallel processes produce fragmented demand signals for industry and limit the government’s ability to plan strategically across the full spectrum of defense production.

This bifurcated system reflects historical conditions in which allied procurement played a relatively smaller role in the overall defense industrial ecosystem. Today, however, allied demand represents a growing share of total procurement for many defense systems. In recent years, FMS have expanded significantly, reaching tens of billions of dollars annually and increasingly shaping production decisions for major defense firms.

At the same time, policymakers and analysts increasingly warn that the U.S. defense industrial base lacks the capacity and resilience required to sustain large-scale military operations while also supplying allies. Recent assessments conclude that the industrial base is struggling to meet the combined demands of U.S. modernization programs and rising allied



procurement requirements in Europe, the Middle East, and the Indo-Pacific (Wasser & Sheers, 2025).

These trends highlight a structural tension within the current acquisition system. While the United States increasingly relies on allies and partners to strengthen collective deterrence, the institutional processes governing defense production continue to treat allied demand as an external variable rather than an integral component of industrial planning.

The consequences of this fragmented demand structure are significant. Production lines frequently experience unstable demand cycles, suppliers face uncertain investment signals, and allied forces encounter lengthy delivery timelines for critical capabilities. These dynamics weaken industrial resilience, raise procurement costs, and slow the fielding of military capability across allied coalitions.

Recognizing these challenges, recent U.S. policy initiatives have begun to more closely link defense export reform with broader industrial strategy. A recent executive order on foreign defense sales reform explicitly states that a more efficient export system can “revitalize the defense industrial base” while simultaneously strengthening allied military capability (White House, 2025).

This paper argues that the United States must shift its conceptual approach to defense demand management. Rather than treating FMS as an external activity managed separately from domestic acquisition, policymakers should integrate allied procurement into a unified Total Demand framework. Under this model, domestic and allied demand would be aggregated into a single planning construct that informs production planning, industrial investment, and long-term sustainment strategies.

Recent institutional developments create an opportunity for such a shift. The alignment of the Defense Security Cooperation Agency (DSCA) under the Office of the Under Secretary of Defense for Acquisition and Sustainment (A&S) creates an opportunity to connect export policy with acquisition planning and industrial strategy. (DoW, 2025; Duffey, 2026). There is a clear recognition of the need to improve visibility of allied and partner demand in order to inform production planning, contracting strategies, and industrial base investments (DoD, 2023).

At the same time rapid advances in artificial intelligence (AI) and data analytics can enable this transformation. Modern analytic tools can synthesize information from previously disconnected datasets including procurement projections, allied requests, readiness data, and industrial capacity information to generate predictive demand forecasts.

This paper proceeds in three sections. First, it examines the structural limitations of the current system, in which allied procurement is treated as an external disruption to domestic acquisition planning. Second, it introduces the concept of a Total Demand framework and outlines how integrated data analytics could support more strategic industrial planning. Third, it explores the broader strategic implications of integrated demand management for industrial resilience, alliance relationships, and coalition interoperability.

## **The Flawed Paradigm: Foreign Military Sales as an Exogenous Disruption**

The contemporary defense acquisition system reflects a historical separation between domestic procurement planning and international defense sales. This institutional divide originates in distinct legal authorities, budgeting processes, and organizational responsibilities that evolved over decades (Congressional Research Service, 2023).

Domestic acquisition within the Department of Defense operates primarily through the Planning, Programming, Budgeting, and Execution (PPBE) system. Through this process, military services develop multiyear procurement plans based on strategic guidance, force



structure requirements, and fiscal constraints. These plans provide industry with relatively predictable demand signals, enabling companies to plan production capacity, workforce levels, and supply chain relationships over extended time horizons.

FMS operates through a fundamentally different process. Allied governments initiate requests for defense articles or services through formal Letters of Request (LOR), which the United States evaluates and processes through the FMS system. Once approved, these requests result in individual FMS cases that are executed alongside, but not fully integrated with, domestic procurement programs. LORs typically arrive after a partner nation has already completed internal budgeting, approvals, and capability decisions (DSCU, 2024). By the time an LOR is submitted the demand signal reaches U.S. industry too late to shape production capacity. Furthermore, they represent a specific purchase request, not the full anticipated life cycle demand from the partner.

While LORs are the formal mechanism that initiates FMS transactions, they provide a late and incomplete signal of partner demand, limiting the ability of the Department of War (DoW) and the defense industrial base to anticipate aggregate allied requirements and plan production capacity accordingly. As a result, the U.S. defense industry rarely receives a comprehensive picture of total global demand for a system (DoD, 2024; National Defense Industrial Association, 2023).

This structural separation produces several challenges for the defense industrial base. One of the most visible effects is the recurring problem of “cold production lines.” When U.S. procurement of a particular system declines, or concludes, production capacity often contracts or shuts down entirely. Years later, when allied demand emerges for the same system, restarting production can require significant time and investment.

These production disruptions occur within a broader industrial ecosystem that analysts already view as fragile. The National Defense Industrial Association’s *Vital Signs* report warns that the United States currently lacks sufficient surge capacity to sustain prolonged conflict while maintaining modernization programs (National Defense Industrial Association [NDIA], 2023).

Fragmented demand signals also complicate long-term management of component obsolescence. Many defense systems rely on specialized components produced by small suppliers that depend on stable production volumes. Without predictable demand signals, these suppliers often exit the market, forcing costly redesigns or delaying sustainment support.

The separation between domestic acquisition and FMS processes contributes to supply chain instability by creating unpredictable and fragmented demand signals, which in turn discourage suppliers from investing in capacity and maintaining inventories, undermining the defense industrial base’s ability to sustain capabilities across U.S. and allied fleets (Defense Business Board, 2025). Without mechanisms to aggregate and integrate domestic and allied demand signals, production capacity expands and contracts episodically rather than being managed as a strategic industrial asset.

In effect, the current system treats the defense industrial base as a reactive supplier responding to episodic procurement decisions rather than as a strategic partner in long-term capability development.

## **A Total Demand Framework**

Addressing these challenges requires a shift in how policymakers conceptualize defense demand management. Rather than viewing allied procurement as an external activity separate from domestic acquisition planning, the United States should treat domestic and allied demand as components of a shared industrial ecosystem. The Total Demand framework proposed in this



paper reflects this conceptual shift. Under this model, demand from U.S. military services and allied partners would be aggregated into a single analytical construct used to inform production planning and industrial investment decisions.

Central to this approach is the integration of multiple data sources that currently exist in institutional silos. Domestic procurement plans including Program Objective Memorandum projections provide one important component of the demand picture. Allied demand signals can be drawn from formal LORs, security cooperation planning documents, and broader defense modernization strategies among partner nations.

The concept aligns closely with emerging priorities within the Department of Defense's National Defense Industrial Strategy, which emphasizes the need for improved production capacity, supply chain resilience, and greater cooperation with allies and partners in defense manufacturing ecosystems (DoD, 2024) as well as priorities reflected by the House Foreign Affairs Committee's bipartisan Foreign Arms Sales Task Force (HFAC, 2025).

Advances in AI and machine learning make it increasingly feasible to synthesize these diverse data streams into predictive demand models. By identifying patterns across historical procurement behavior, geopolitical events, and readiness indicators, analytic systems could forecast likely demand trajectories years in advance.

### **AI-Powered Data Integration: Creating the "Total Demand" Signal**

The placement of DSCA within A&S creates the organizational authority to demand data sharing. AI-powered data analytics creates technical capability to ingest, harmonize, and analyze that data from its disparate sources (Gahane et al., 2025).

An AI-powered system could be used break down the traditional silos which have characterized domestic procurement and FMS processes to create a unified demand signal. The operational benefits of this approach are substantial. Integrated demand forecasts could support minimum sustaining production rates, enable larger block buys, and stabilize supply chains for critical components. These mechanisms would strengthen industrial resilience while lowering procurement costs for both the United States and its partners.

Defense Logistics Agency (DLA) has highlighted the potential for AI to assist with mitigating Supply Chain Risk Management by improving visibility and forecasting customer demand to optimize supply chain workflows, and recommending alternative, pre-qualified suppliers during disruptions (DLA, 2025). These systems analyze historical demand, real-time logistics data, and supplier performance to detect patterns and anticipate future needs, improving forecasting accuracy and enabling more proactive planning across the defense industrial base.



**Table 1. Data Sources to Feed a Total Demand Signal**

| <b>Data Source</b>                      | <b>Traditional Silo</b>                                                                                                                      | <b>AI-Powered Harmonization &amp; Insight</b>                                                                                                                                                                                   |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DoW Program Objective Memorandums (POM) | Lives within the Pentagon's programming and budgeting world. Seen by the military services and program offices.                              | AI ingests the POM data, mapping out the DoD's planned procurement quantities and timelines for systems like the F-35 or Standard Missile-6 over the next 5–10 years.                                                           |
| Partner Nation LORs                     | Handled by DSCA, the State Department, and regional Combatant Commands. Often seen as a "transaction" separate from DoD long-range planning. | AI ingests all incoming LORs, digitizing and categorizing them by system, quantity, and requested delivery date.                                                                                                                |
| NLP-Derived Early Signals               | Scattered across diplomatic cables, intel reports, and public statements. Often just "background noise."                                     | AI processes this unstructured data to generate a <i>probabilistic forecast</i> of future LORs, complete with confidence scores (e.g., "75% probability of a request for 24 HIMARS launchers from Country Y within 18 months"). |
| Allied Readiness Data                   | Held by partner nations or within U.S. security cooperation offices. Often in non-standard formats (spreadsheets, emails).                   | AI ingests readiness reports on existing allied systems, predicting future demand for spare parts and replacements across the entire allied fleet.                                                                              |
| Industrial Base & Supply Chain Data     | Resides with prime contractors and their sub-tier suppliers. Often proprietary and not shared.                                               | The AI platform can receive anonymized data from industry on production capacity, lead times for critical subcomponents (like rocket motors or microelectronics), and current backlogs.                                         |

By harmonizing these diverse data sources an AI platform designed to integrate foreign and domestic demand could create a comprehensive, real-time dashboard of total global demand.

- Solving the "Cold Production Line" Problem:
  - Old Way: The U.S. Army buys its last PATRIOT missile in 2026. The line shuts down. In 2027, three allies submit LORs. A&S and industry scramble to restart the line at great expense.
  - New Way: The AI model, analyzing DoD POM data and NLP-derived forecasts, projects a dip in domestic demand in 2026 but sees a high probability of 3–5 partner orders in 2027–2028. A&S can now use this integrated demand signal to authorize a "minimum sustaining rate" for the PATRIOT production line, funded by a combination of DoD and anticipated FMS funds, keeping it "warm" and ready. The partner's order is filled faster and at a lower cost.
- Achieving Economies of Scale:
  - Old Way: The U.S. Navy orders 100 Naval Strike Missiles. Six months later, a partner orders 20. A third partner orders 30 more a year later. Each is treated as a separate order with its own contract and pricing.





- **New Way:** The AI aggregates the Navy's demand, the first partner's firm LOR, and a high-confidence forecast for the third partner's future request. A&S can now approach industry with a single "block buy" request for 150+ missiles. This larger, more predictable order allows the contractor to secure better prices on raw materials and optimize production, lowering the unit cost for *both* the U.S. Navy and its allies.

## **Natural Language Processing as an Early Warning System for Future Defense Requirements**

A core challenge in the efforts to improve FMS efficiency is that formal LORs occur only at the end of a lengthy process of internal deliberation, budgeting, and political decision-making within partner nations (Chindea, 2024). Relying solely on LORs as a demand signal leaves the U.S. defense industrial base (DIB) perpetually in a reactive posture, unable to anticipate partner requirements in advance. Natural Language Processing (NLP) offers a potential solution by enabling the DoW to analyze the vast stream of pre-decisional communications, reports, and open-source publications that precede formal requests (Schirmer, 2024). These unstructured sources can provide early, predictive signals of emerging demand, allowing planners and industry to anticipate requirements before LORs are officially submitted.

This includes:

- **Diplomatic Traffic:** Reports and summaries from embassies and military attachés.
- **Partner Nation Media:** News articles, press releases, and televised interviews with defense ministers or military chiefs.
- **Official Publications:** Defense white papers, national security strategies, and parliamentary records from allied governments.
- **Intelligence Reporting:** Summaries and analyses regarding regional threats and capability gaps.

Manually reading and synthesizing this volume of information is impossible for any human team. NLP algorithms, however, can process millions of documents in near-real-time to identify patterns and signals that would otherwise be missed. By applying NLP to unstructured data associated with FMS can generate early warning signals of emerging capability requirements, anticipate allied demand, and integrate these insights into a predictive Total Demand framework, reducing production lag, stabilizing the industrial base, and enhancing coalition readiness (DoD, 2022).

NLP offers important potential in efforts to reform FMS to improve efficiency and ensure that it becomes more anticipatory. The following are a few examples of specific NLP techniques that could assist in better exploiting unstructured data relevant international defense markets to create better demand signals for U.S. industry (Seowet al., 2025).



**Table 2. Key NLP Techniques and Their Applications**

| NLP Technique                  | Application in FMS Context                                                                                                                                                                                                   | Example in Practice                                                                                                                                                                                                                                                                    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Named Entity Recognition (NER) | NER acts like a high-speed scanner, automatically identifying and categorizing key information. It finds mention of specific weapon systems, military units, government officials, countries, and private defense companies. | An NLP model scans thousands of articles from a partner nation's media. It tags every mention of "F-16," "Patriot missile," "maritime patrol," and "cybersecurity," creating a structured database of who is talking about what.                                                       |
| Topic Modeling                 | This technique automatically discovers the hidden themes or "topics" in large volumes of text. It can detect a subtle but persistent shift in a country's defense priorities over time.                                      | By analyzing a partner's parliamentary debates over 12 months, the model might detect that the frequency of the "arctic surveillance" and "icebreaker capability" topic cluster is steadily increasing, signaling a new, emerging priority long before a specific vessel is requested. |
| Sentiment Analysis             | Sentiment analysis assesses the emotional tone of the text, determining if it's positive, negative, or neutral. This helps gauge urgency, concern, or commitment related to a specific defense issue.                        | A foreign defense minister gives a speech about the country's aging helicopter fleet. Sentiment analysis can score the urgency and concern in their language, indicating how high a political priority its replacement has become.                                                     |
| Relationship Extraction        | This technique identifies the relationships between different entities. It doesn't just see a mention of a weapon system; it understands the context connecting it to a threat, a location, or a capability gap.             | The algorithm processes an intelligence report and extracts a key relationship: "[Country X's] Navy <i>lacks capability against</i> [adversary's new submarine class], <i>creating a requirement for</i> [anti-submarine warfare aircraft]."                                           |

By combining these techniques, the DoW can move from disconnected pieces of information to a cohesive, predictive intelligence picture. Instead of just knowing a country is *interested* in air defense, analysts can see:

- Which specific systems are being discussed most often (NER).
- How the priority of air defense compares to other needs, like coastal security (Topic Modeling).
- How urgently political leaders are framing the need (Sentiment Analysis).
- What specific threat is driving the requirement (Relationship Extraction).

This provides a rich, multi-dimensional early signal of future demand, allowing the DoW to begin conversations with industry, assess supply chain readiness, and streamline the eventual FMS case long before a formal request is ever signed

## Strategic Implications of a Total Demand Framework

Beyond its immediate industrial and acquisition benefits, a Total Demand framework carries significant strategic implications for the United States and its allies. Integrating domestic procurement planning with allied demand signals has the potential to strengthen the resilience



of the defense industrial base, reinforce alliances, and enhance coalition warfighting capability in an era characterized by intensifying strategic competition and growing reliance on multinational military operations.

## **Industrial Resilience**

First, integrating domestic and allied demand signals would strengthen the resilience of the U.S. defense industrial base. Stable and predictable demand is one of the most important factors shaping industrial investment decisions. When production orders fluctuate dramatically, firms and suppliers are reluctant to invest in workforce development, manufacturing capacity, or supply chain expansion. Conversely, more stable demand signals encourage sustained investment and allow companies to maintain skilled labor forces, specialized tooling, and production infrastructure.

A Total Demand approach would provide industry with a more comprehensive picture of global procurement requirements for key defense systems. By aggregating U.S. military requirements with anticipated allied purchases, policymakers could establish more stable production baselines and maintain minimum sustaining production rates for critical systems. This stability reduces the risk of supplier attrition, which has become a growing concern within the defense industrial ecosystem. Recent assessments of the industrial base highlight that unpredictable procurement cycles contribute to supplier exits, workforce instability, and production bottlenecks across key sectors (DoD, 2024; NDIA, 2023).

Regional efforts like the Partnership for Indo-Pacific Industrial Resilience (PIPIR) highlight how integral allies are to achieving industrial resilience. PIPIR is a multilateral forum of 14 Indo-Pacific and Euro-Atlantic partners collaborating to accelerate Indo-Pacific contributions to global defense industrial base resilience. Working with industry, capital providers, and non-government stakeholders, PIPIR is creating an “ecosystem of information exchange, technical cooperation, supply chain resilience, and co-production and co-sustainment collaboration” (DoW, 2025).

Jedidiah Royal, principal deputy assistant secretary of war for Indo-Pacific security affairs, has described these efforts as “the first bricks being laid in the foundation of a stronger, more agile, more resilient defense industrial base across the region” (Luckenbaugh, 2026).

Strengthening industrial resilience is especially important in the context of potential high-intensity conflict. Large-scale operations would require the rapid expansion of production across multiple defense sectors, including precision munitions, missile defense interceptors, and advanced electronic systems. Without stable peacetime production lines and a healthy supplier network, such surge capacity would be difficult to achieve. A Total Demand framework therefore helps ensure that the industrial base remains capable of supporting both routine procurement and potential wartime expansion.

## **Strengthening Alliances**

Second, integrating allied demand into acquisition planning would enhance the United States’ ability to support allies and partners, reinforcing alliances as instruments of strategic statecraft. Security cooperation and defense exports play a critical role in enabling partner nations to develop military capabilities aligned with U.S. operational concepts and technologies. When these systems are delivered efficiently and affordably, they strengthen the credibility of U.S. security commitments and deepen long-term defense relationships.

However, delays in defense exports have historically frustrated partner nations and undermined the strategic benefits of security cooperation. In January of 2026, Japan’s Board of Audit revealed that military equipment worth roughly 1.1 trillion yen (\$7.1 billion) procured from the United States had yet to be delivered more than five years after contracts were signed.



These delays became a focus of bilateral talks later that month (Manuel, 2026). Lengthy acquisition timelines, supply chain disruptions, and production bottlenecks can delay delivery of critical capabilities for years. These delays may encourage partner nations to seek alternative suppliers, including purchasing from competitors whose defense export strategies often emphasize speed and flexibility (Winn, 2025).

A Total Demand framework can help mitigate these challenges by enabling more proactive planning for allied demand. By incorporating allied procurement projections into production planning, the United States could align domestic and partner acquisitions in ways that reduce lead times and lower unit costs through larger procurement quantities. These efficiencies make U.S. systems more attractive to partner nations while improving the credibility of American security cooperation commitments.

More broadly, the ability to deliver advanced capabilities quickly and reliably strengthens the strategic position of the United States within global defense markets. Defense exports are not purely economic transactions; they also shape long-term military relationships, interoperability, and geopolitical alignment. By improving the efficiency and predictability of the FMS system, a Total Demand approach enhances the role of security cooperation as a strategic instrument for strengthening alliances and countering the influence of competing defense suppliers (Congressional Research Service, 2023).

### **Forging a More Lethal Coalition**

Finally, a Total Demand framework has important implications for coalition warfighting capability. Modern military operations increasingly rely on multinational coalitions composed of allied forces operating across multiple domains. The effectiveness of these coalitions depends heavily on interoperability—the ability of participating forces to share information, operate compatible equipment, and sustain operations through integrated logistics networks.

Synchronizing procurement and modernization cycles across allied forces represents a key step toward achieving such interoperability. When nations acquire and upgrade systems independently, their capabilities often diverge over time. Differences in software configurations, sensor packages, communications systems, or sustainment arrangements can complicate coalition operations and limit the ability of forces to integrate effectively in combat environments.

By aggregating demand signals across allied and U.S. forces, a Total Demand framework could facilitate greater synchronization in the procurement and modernization of shared systems. Coordinated acquisition schedules make it easier to align software upgrades, configuration updates, and sustainment planning across fleets operated by multiple nations. Over time, this alignment enables the development of integrated operational concepts and shared logistical infrastructures. In this sense, a Total Demand framework contributes not only to industrial efficiency but also to coalition lethality. By enabling more integrated capability development among allies, the United States can help build coalitions that are more interoperable, more responsive, and more capable of deterring or defeating potential adversaries.

FMS is not an exogenous “foreign” affair that happens to the DoW acquisition system. It is an integral, vital component of the acquisition system and a primary tool for managing the health of the defense industrial base. The new, integrated paradigm views FMS demand as another critical data stream that must be factored into the overall U.S. acquisition strategy. It sees partner demand not as a disruption, but as an opportunity for stabilization and efficiency.



**Table 3. Conceptual Shifts of the Total Demand Framework**

| <b>Aspect</b>      | <b>Old Paradigm (FMS as Exogenous)</b>                                       | <b>New Paradigm (FMS as a Key Component)</b>                                                                               |
|--------------------|------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| Core Identity      | FMS is a foreign policy transaction.                                         | FMS is an industrial base strategic tool.                                                                                  |
| Demand Signal      | Unpredictable, reactive “orders” from outside.                               | A predictable, forecastable stream of “total global demand.”                                                               |
| Impact on Industry | A source of disruption (e.g., cold line restarts) and unpredictable revenue. | A source of stabilization (e.g., production smoothing, warm lines) and predictable, recurring revenue.                     |
| Primary Goal       | Fulfill a specific partner request.                                          | Shape and sustain the industrial base to support a ready and interoperable U.S. and Allied coalition.                      |
| Stakeholder Focus  | Primarily focused on the needs of the foreign partner as a “customer.”       | Focused on the dual benefit to the partner (capability) and the United States (industrial base health & interoperability). |

### **New Incentives for Allied Industrial Participation**

A Total Demand framework may also reshape the role of allied defense industries within the broader coalition defense ecosystem. When procurement planning accounts for aggregated demand across U.S. and partner forces, allied governments gain stronger incentives to participate directly in the production and sustainment of shared defense systems. Rather than functioning primarily as end users of U.S.-produced equipment through FMS transactions, partner nations may seek deeper involvement in the industrial processes that support these capabilities.

Such participation can take several forms. Governments may negotiate licensed production arrangements, allowing domestic industries to assemble or produce elements of U.S.-designed systems under agreed export control frameworks. In some cases, allies may pursue co-development programs, contributing financial resources, technical expertise, or manufacturing infrastructure to jointly develop next-generation capabilities. Finally, allies may play a growing role in sustainment supply chains, including maintenance, repair, and overhaul (MRO) activities that support shared fleets operating across multiple nations.

A Total Demand framework strengthens the logic for these arrangements because it emphasizes the scale and durability of coalition demand. When industry can anticipate long-term production volumes driven by both U.S. and allied procurement, it becomes more feasible to distribute portions of manufacturing and sustainment activities across trusted partner nations. Emerging multinational initiatives also aim to operationalize this concept by developing distributed industrial networks among allies. Programs such as the PIPIR seek to coordinate production, sustainment, and supply chains across partner nations to strengthen the resilience of coalition defense industrial capacity. This distributed model can expand overall production capacity while reducing reliance on single-country supply chains that may prove vulnerable during periods of crisis or conflict.

The development of distributed allied manufacturing networks is increasingly viewed by defense planners as an important component of industrial resilience. Diversifying production and sustainment activities across allied economies can mitigate these risks while also enabling coalition partners to contribute more directly to shared defense capabilities (Aldisert & Cook,





2025). Recent RAND research emphasizes that modern defense production increasingly depends on cooperation among allied industrial bases. Because no single nation possesses sufficient surge capacity across all sectors, coalition operations require coordinated production and supply networks that span multiple countries. Strengthening interoperability among allied defense industries can therefore expand production capacity, mitigate supply chain vulnerabilities, and improve the ability of coalition forces to sustain operations in prolonged conflicts (RAND, 2024).

Beyond resilience, deeper industrial integration may also reinforce the political foundations of defense cooperation. When allied industries are directly involved in producing or sustaining shared military capabilities, governments often develop stronger long-term commitments to those systems and to the partnerships that sustain them. This dynamic has been observed in several multinational defense programs where industrial participation helped solidify enduring security relationships. For example, research on multinational fighter aircraft programs finds that industrial workshare and co-production arrangements help sustain partner commitment to joint programs by aligning economic and strategic interests (Lorell et al., 2013).

In this sense, a Total Demand framework has the potential to expand the role of defense exports beyond transactions between governments. By encouraging collaborative production and sustainment arrangements, it could foster a more integrated coalition defense industrial ecosystem in which allied economies contribute to the development, production, and long-term sustainment of shared military capabilities. Such integration would not only improve industrial resilience but also strengthen the strategic cohesion of alliances operating in an increasingly contested global security environment.

## Conclusion

The United States faces a strategic environment in which military advantage increasingly depends on the ability to produce, sustain, and field capabilities at scale across a coalition of partners. In this context, the traditional separation between domestic defense acquisition and FMS is no longer sustainable. A system designed in an earlier era when defense exports were viewed primarily as diplomatic transactions rather than components of industrial strategy now imposes significant inefficiencies on the defense industrial base and slows the delivery of critical capabilities to both U.S. and allied forces.

By reconceptualizing FMS as a central component of a unified demand management strategy and by integrating domestic acquisition planning with allied procurement requirements, the DoW can adopt a framework that reflects the full scope of global demand for key defense systems. Such a framework would transform both how the United States manages its industrial base as well as how it delivers capabilities to allies and partners.

Advances in data analytics and AI provide the tools necessary to operationalize this shift. By aggregating data from sources such as POMs, LORs, security cooperation engagement data, industrial capacity metrics, and partner readiness indicators, AI-enabled systems could generate predictive forecasts of global demand for defense capabilities by capturing emerging signals in unstructured information sources well before formal procurement requests are submitted. The result would be a real-time understanding of coalition demand that allows policymakers and industry to anticipate future requirements rather than react to them.

The strategic implications of such a system extend beyond efficiency. A Total Demand approach could significantly strengthen the resilience of the defense industrial base by stabilizing. Integrating allied demand into procurement planning would encourage deeper industrial participation among partner nations. Allies that anticipate sustained demand for shared systems may seek greater involvement in component manufacturing, licensed



production, or sustainment supply chains. Industrial policy should more actively incentivize allied participation in production and sustainment, not just procurement. By expanding licensed production, co-production, and component sourcing among trusted partners, the United States can build a more distributed defense industrial base that increases surge capacity and reduces supply chain risk in a crisis. FMS can be used as the entry point for this integration: initial sales create a common platform, which can then evolve into shared sustainment networks, regional maintenance hubs, and eventually co-development opportunities.

Over time, this dynamic could support the development of distributed allied defense industrial networks capable of expanding production capacity and mitigating supply chain disruptions in times of crisis. In this sense, Total Demand offers an opportunity to align procurement planning, industrial policy, and security cooperation to foster an integrated defense industrial ecosystem that spans allied economies. Such a system would not only enable faster and more affordable capability delivery but also reinforce the political and military cohesion of allied partnerships.

Realizing this vision will require deliberate action. Legislative reforms must create the authorities necessary to support flexible procurement strategies and industrial base investments. Departmental policies must establish common data architectures and encourage responsible data sharing between government and industry. Just as importantly, the defense acquisition and security cooperation workforce must adopt a mindset that views FMS not as discrete transactions but as instruments of strategic industrial management and coalition capability development.

The realignment of the DSCA under A&S offers a unique opportunity to advance this transformation. If paired with the adoption of a Total Demand framework supported by modern data analytics and AI, this organizational change could help create a more resilient and responsive defense industrial base while improving responsiveness to ally and partner demand.

## References

- Aldisert, A., & Cook, C. R. (2025). *Improving arms sales, technology transfer, and defense industrial cooperation with allies and partners*. Center for Strategic and International Studies.
- Chindea, I. A., Moroney, J. D. P., Webber, S., Brin, I. M., Thaler, D. E., Rhoades, A. L., Adler, A., Grill, B., Cormarie, P., Lashendock, J., et al. (2024, July 24). *Aligning strategic priorities and foreign military sales to fill critical capability gaps*. RAND Corporation. [https://www.rand.org/pubs/research\\_reports/RRA2438-2.html](https://www.rand.org/pubs/research_reports/RRA2438-2.html)
- Congressional Research Service. (2023). *U.S. arms sales: Background and issues for Congress* (CRS Report R46384).
- Defense Business Board. (2025). *Supply chain illumination report* (Report FY25-02).
- Defense Logistics Agency. (2024). *Utilization of artificial intelligence (AI) to illuminate supply chain risk*.
- Defense Security Cooperation University. (2024).
- DoD. (2024). *National defense industrial strategy*.
- DoW. (2025, June 1). *Partnership for Indo-Pacific industrial resilience (PIPIR)* [Fact sheet]. <https://media.defense.gov/2025/Jun/02/2003730341/-1/-1/1/FACT-SHEET-PARTNERSHIP-FOR-INDO-PACIFIC-INDUSTRIAL-RESILIENCE.PDF>



- DoW. (2025, November 7). *Unifying the department's arms transfer and security cooperation enterprise to improve efficiency and enable burden-sharing* [Memorandum].
- Duffey, M. P. (2026, March 17). *Prepared remarks before the U.S. House of Representatives Committee on Foreign Affairs: Reforming America's defense sales* (119th Cong., 2nd sess.).
- Gahane, S., Dubey, A., Anawade, P., & Sharma, D. (2025). The use of artificial intelligence and predictive data analytics approaches and techniques in AI-driven modern applications and sectors. In M. Tuba, S. Akashe, & A. Joshi (Eds.), *ICT systems and sustainability. ICT4SD 2024. Lecture notes in networks and systems* (Vol. 1194, pp. 211–223). Springer. [https://doi.org/10.1007/978-981-97-9523-9\\_18](https://doi.org/10.1007/978-981-97-9523-9_18)
- House Foreign Affairs Committee. (2025, September 2). *House passes HFAC-led reforms to foreign arms sales process* [Press release]. <https://foreignaffairs.house.gov/media/press-releases>
- Lorell, M., et al. (2013). *Multinational fighter programs: Lessons learned*. RAND Corporation.
- Luckenbaugh, J. (2026, March 9). *U.S., allies working to solidify 'wobbly' Indo-Pacific defense industrial base*. National Defense Magazine. <https://www.nationaldefensemagazine.org/articles/2026/3/9/us-allies-working-to-solidify-wobbly-indo-pacific-defense-industrial-base>
- Manuel, R. (2026, January 26). *Japan audit flags \$7.1B in undelivered U.S. military equipment after five years*. The Defense Post. <https://thedefensepost.com/2026/01/26/japan-undelivered-us-equipment/>
- National Defense Industrial Association. (2023). *Vital signs 2023: The health and readiness of the defense industrial base*.
- RAND Corporation. (2024). *Implications of the Russia–Ukraine war for the U.S. and allied defense industrial bases*.
- Schirmer, P., Jaycocks, A., Mann, S., Marcellino, W., Matthews, L. J., Parsons, J. D., & Schulker, D. (2021, July 28). *Natural language processing: Security- and defense-related lessons learned*. RAND Corporation. <https://www.rand.org/pubs/perspectives/PEA926-1.html>
- Seow, W. L., Chaturvedi, I., Hogarth, A., et al. (2025). A review of named entity recognition: From learning methods to modelling paradigms and tasks. *Artificial Intelligence Review*, 58, 315. <https://doi.org/10.1007/s10462-025-11321-8>
- Wasser, B., & Sheers, P. (2025). *From production lines to front lines: Revitalizing the U.S. defense industrial base for future great power conflict*. Center for a New American Security.
- White House. (2025). *Executive order: Reforming foreign defense sales to improve speed and accountability*.
- Winn, A., McCulloch, C., Moroney, J. D. P., Pezard, S., Thaler, D. E., Grill, B., Marcinek, K., & Kelly, L. (2025, April 17). *U.S. options for identifying third-party suppliers to meet ally and partner capability needs*. RAND Corporation. [https://www.rand.org/pubs/expert\\_insights/PEA2822-1.html](https://www.rand.org/pubs/expert_insights/PEA2822-1.html)



# Embracing Non-Traditional Solutions: Evolving Security Cooperation for the Modern Threat Environment

**Jarrett M. Lane**—is an Adjunct Research Staff Member at the Institute for Defense Analyses. His work primarily focuses on security cooperation and innovation. In addition to his role at IDA, Jarrett is Founder and Chair of the North Carolina Critical Technologies Alliance, a non-profit focused on promoting growth and development of North Carolina's defense and "deep tech" industries. He previously worked in government affairs and business development at Amazon and AWS. He also served as Chief of Staff at the Section 809 Panel, a defense acquisition reform initiative mandated by Congress. [jlane@ida.org]

**Jennifer M. Taylor**—is a Research Staff Member at the Institute for Defense Analyses currently serving as co-lead for IDA's security cooperation, institutional capacity building, and allies and partners work. Jennifer has more than 25 years' experience in government, think tanks, and as a consultant, transforming government institutions through strategy, planning, analysis, and implementation. She has a personal commitment to building better government institutions through making connections that were hidden from view. [jtaylor@ida.org]

## Abstract

Advancing burden-sharing and enhancing the resilience among partners requires the security cooperation enterprise to open the aperture to a wider array of technologies and solutions. The security cooperation enterprise's near-exclusive reliance on "traditional" solutions covered by programs of record to build partner nation capabilities under-values the criticality of commercial technologies in modern warfighting. Further, neo-primes and early-stage defense technology companies are rapidly developing new military-unique technologies that may not become a program of record but could prove useful to allies and partners.

The increasingly perilous threat environment, compounded by the pace of technological change, means the security cooperation enterprise must accelerate efforts to help partners acquire non-program of record solutions. As the security cooperation enterprise undergoes significant changes (i.e., Defense Security Cooperation Agency's realignment to the Undersecretary of War for Acquisition & Sustainment) now is an opportune time to evolve the Department's approach to security cooperation to leverage these defense technology opportunities.

There are some practical steps that the Department could take immediately. For example, organizations steeped in commercial technology and innovation (e.g., the Defense Innovation Unit) should be formally incorporated into the security cooperation enterprise and resourced to support security cooperation initiatives. However, bolder steps to create entirely new acquisition pathways for allies and partners should also be considered.

Authorities could be revised to allow foreign military financing to be used by all allies and partners to procure certain technologies through direct commercial contracts. Marketplaces of vetted, trusted technologies would allow allies and partners to more quickly acquire new capabilities. Adopting a portfolio-like approaches to aligning partner nation requirements with technical expertise resident in the Foreign Area Officer community could also enable security cooperation organizations to more effectively identify and implement solutions for partner nations.

## Background

In fiscal year 2024 alone, the United States transferred \$117.9 billion worth of defense articles to allies and partners (Department of State, 2025). The vast majority of what is provided to allies and partners are linked to programs of record and traditional capabilities, meaning the United States military's procurement of those same products and services is baked into the Department of Defense's (DoD) budget, funded by Congress, and will therefore be supported by the defense industry long-term.



Focusing arms transfers on technologies and products covered by programs of record has four distinct advantages. First, technologies covered by programs of record are typically tested rigorously to ensure they meet the U.S. military's own standards and requirements for safety, reliability, operational efficacy, and more. This level of rigor inspires confidence among allies and partners that they are acquiring a high-quality product.

Second, transferring the same technologies developed for the U.S. military to allies and partners helps promote interoperability. Though interoperability with allies and partners is hard to achieve, it remains a goal of the DoD. Common understanding of systems and shared architectures can enable the United States and its allies and partners to fight as a combined force.

Third, because technologies covered by programs of record are typically acquired by the United States for decades, allies and partners that acquire those same technologies are effectively investing in a long-term relationship with the United States. Finally, encouraging allies and partners to acquire technologies covered by programs of record helps bolster the United States' defense industrial base (DIB). Global sales represent an important revenue stream for the DIB and helps ensure the DIB is not exclusively reliant on U.S. military sales to keep production lines running and continue offering associated services (e.g., maintenance, training, etc.).

There are limitations, however, to focusing on technologies covered by programs of record, particularly in the context of helping allies and partners build capabilities and ultimately assume greater responsibilities for burden-sharing. First, the DIB and the foreign military sales (FMS) process are unable to deliver solutions quickly enough to allies and partners. Right now, there is a \$250 billion FMS backlog (Zinke, 2026). Though delays cannot be exclusively attributed to the United States—allies and partners have a role in executing sales agreements, too (Gomez, 2024)—delays nonetheless present an opening to adversaries and competitors. Failure to deliver timely, cost-effective solutions to solve allies' and partners' operational problems may lead them to explore offerings from other countries, including China and Russia (Kurilla, 2025).

Second, technologies covered by programs of record may not be best-suited to address new challenges. In the course of the authors' past research and work on the security cooperation enterprise, U.S. military groups (USMILGRPs) highlighted difficulties responding quickly to emergent allied and partner requirements, such as small unmanned aerial systems (sUAS) and counter-drone solutions.

The security cooperation enterprise's difficulties responding to emergent requirements speaks to a broader challenge shared by the U.S. military and its allies and partners: keeping pace with the speed of technological change and adapting to the trajectory of future warfare. Further opening the aperture to non-program of record solutions, including commercially-available technologies, will be critical to helping allies and partners build capabilities that are resilient, sustainable, and effective in the modern environment.

This paper builds upon research conducted by the Institute for Defense Analyses (IDA), which was sponsored by the Defense Security Cooperation Agency (DSCA) in 2024, that explored ways to more systematically and holistically incorporate commercial technologies into security cooperation efforts. This paper broadens the aperture to consider non-program of record solutions overall, including defense-specific capabilities that are not covered by programs of record, and explores opportunities and challenges at the intersection of security cooperation, DoD acquisition, and innovation.





This paper starts by providing an overview of major recent developments in security cooperation policy, particularly at the intersection of defense acquisitions and innovation. It then provides insights into the importance of commercial technologies in modern warfighting and some of the challenges the security cooperation enterprise faces in systematically incorporating commercial technologies into how it builds allied and partner capabilities. The paper concludes with recommendations for the Department and Congress to consider for ensuring the United States is able to offer optimal solutions to allies and partners in a timely, cost-competitive manner.

## Recent Policies on FMS

Significant changes are occurring in both the defense acquisition system, the defense industrial base, and in security cooperation. Discerning the opportunities and challenges that these changes present (particularly with respect to allies, partners, and advancing burden sharing) first requires an understanding of major policy changes and initiatives with respect to each. This section provides a timeline of recent major efforts and a brief synopsis of their implications.

- **August 2022: DoD FMS Tiger Team.** Secretary of Defense Lloyd Austin establishes a tiger team to address “inefficiencies in the United States’ transfer of defense articles and services” to allies and partners. The tiger team is led by the Deputy Undersecretary of Defense for Acquisition and Sustainment the Deputy Undersecretary of Defense for Policy (DoD, 2023).
- **June 13, 2023: FMS Tiger Team Tasking Memo.** Secretary Austin directs implementation of recommendations delivered by the DoD’s FMS Tiger Team. Key directives include establishing the Defense Security Cooperation Service (DSCS); steps to improve the efficiency of technology releases; *developing a methodology for facilitating non-program of record transfers*; standardizing FMS contract awards and process maps; incorporating allied and partner requirements into DIB production capacity; and working with the Department of State and Congress on FMS process improvements (DoD, 2023).
- **June 27, 2023: House Foreign Affairs Committee FMS TIGER Task Force.** Just days after the DoD released its FMS Tiger Team recommendations, House Foreign Affairs Committee Chairman Michael McCaul announced an FMS Technical, Industrial, and Governmental Engagement for Readiness (TIGER) Task Force (House Committee on Foreign Affairs, 2023). The task force was bipartisan and included representatives on both the House Foreign Affairs and Armed Services committees. The task force was directed to develop legislative fixes (rather than policy and process changes internal to State and the DoD) to the FMS process (House Committee on Foreign Affairs, 2024).
- **December 22, 2023: Fiscal Year 2024 National Defense Authorization Act.** Section 873 gave Combatant Commanders more input and authorities pertaining to FMS, including a directive to the Undersecretary of Defense for Acquisition and Sustainment to work with the Combatant Commanders (among others) to develop a list of systems that would benefit from exportability to support security cooperation objectives. Congress also made Australia and the United Kingdom priorities for expedited FMS case processing and technology transfers.
- **February 7, 2024: FMS TIGER Task Force Report.** The task force’s report points to multiple problems in the FMS system including, but not limited to, thresholds for congressional review, lack of accountability within the DoD or State for delays, the DoD undervaluing the benefit of exportability, and lack of transparency and traceability in the



FMS process. Select recommendations included changing Congressional notification thresholds; leveraging the Special Defense Acquisition Fund (SDAF) to stockpile high-demand items and help overcome inherent challenges forecasting future demand; incentivizing security cooperation organizations (SCOs) to help allies and partners acquire the optimal solution, including non-program of record solutions; and creating processes and mechanisms to pre-clear and expedite non-program of record transfers (House Committee on Foreign Affairs, 2024).

- **December 23, 2024: Fiscal Year 2025 National Defense Authorization Act.** Section 1210 required the Secretary of Defense to establish the FMS Continuous Process Improvement Board (CPIB). The CPIB is tasked with improving understanding of allied and partner requirements, speeding up technology release reviews, and providing allies and partners with “relevant priority equipment” (U.S. Congress, 2024).
- **April 9, 2025: Executive Order 14268, “Reforming Foreign Defense Sales to Improve Speed and Accountability.”** The principal thrust of this executive order (EO) is to make the FMS process faster and, by extension, accelerate growth of the defense industrial base. Consistent with the FMS TIGER Task Force’s recommendations, the EO pushes for greater exportability and requires State and DoD to deliver a proposal to Congress regarding changes to notification thresholds (The White House, 2025).
- **November 7, 2025: Memorandum on “Unifying the Department’s Arms Transfer and Security Cooperation Enterprise.”** Secretary Hegseth directs the realignment of the Defense Security Cooperation Agency (DSCA) and the Defense Technology Security Administration (DTSA) from the Undersecretary of Defense for Policy (USD[P]) to the Undersecretary of Defense for Acquisition and Sustainment (USD[A&S]). Realigning DSCA to USD(A&S) ensures foreign demand for U.S. weapons and platforms is factored into the Department’s own forecasting and that there is greater cohesion in the front-to-end process (i.e., acquisition through delivery) for weapons and platforms purchased by both the United States and foreign buyers (DoD, 2025).
- **December 4, 2025: National Security Strategy (NSS).** The NSS elevates industrial policy as a core national security imperative. Consistent with EO 14268, the NSS connects weapons sales with strengthening supply chains and economic resilience. The NSS also codifies the concept of burden-sharing and burden-shifting, which offers “more favorable treatment” on arms sales and technology sharing to countries that “take more responsibility for security in their neighborhoods” (NSS, 2025).
- **December 18, 2025: Fiscal Year 2026 National Defense Authorization Act.** Arms sales and transfers policy is heavily featured in this NDAA, alongside numerous defense acquisition reform provisions. Notably, Section 1214 requires the Secretary of Defense to establish an office to support sales of non-program of record systems to allies and partners by October 1, 2026. The office is responsible for coordinating with allies and partners on non-program of record purchases; facilitating industry engagements with foreign customers; developing FMS cases for non-program of record systems; and coordinating internally within the DoD to expedite delivery of non-program of record systems.
- **January 23, 2026: National Defense Strategy (NDS).** Burden-sharing is one of the NDS’s four lines of effort; in that vein, the Department will prioritize arms sales and defense industrial collaboration (NDS, 2026).
- **February 6, 2026: EO 14383 “Establishing an America First Arms Transfer Strategy.”** This EO builds upon EO 14268 by clarifying that arms sales will be used to



build production capacity for critical weapons. It also requires a sales catalog of priority systems to be marketed to allies and partners. Critically, the EO highlights how arms transfers can help “new entrants and nontraditional defense companies” innovate, grow, and contribute to the DIB (The White House, 2026).

- **February 10, 2026: DSCA and DTSA Realign to USD(A&S).** The realignment creates an organizational structure designed to implement policy directives to ensure the Department’s own acquisitions and arms sales reinforce growth of the DIB (Defense Security Cooperation Agency, 2026).

The timeline shows that at the policy levels—both Congressional and within the Executive—there is shared recognition that the FMS system needs reforms to deliver optimal solutions (regardless of whether the solution is part of a program of record or not) to allies and partners in a timely manner. Doing so is crucial to both ensuring the United States remains the premier partner of choice for arms sales, as well as advancing the current strategic objective of enhancing burden-sharing among allies and partners.

What is less clear at this time, however, is the breadth of non-program of record solutions that the Department will consider priorities for helping allies and partners acquire. There are infinitely more non-program of record solutions than there are programs of record. Discernment (i.e., limiting the focus to those that can have the highest impact and are most readily employable both allies and partners) will be critical to meeting policymakers’ intent to expedite delivery of optimal solutions to increase burden-sharing.

Further, there is no explicit recognition of the role that commercial-first companies and technologies should play in building allied and partner capabilities. Emphasis on the importance of increasing foreign sales to fuel growth and strengthening of the DIB is clear. Viewed in the context of policy imperatives to grow the DIB, reference to helping “new entrants” sell to allies and partners seems principally targeted at neo-primes and emergent defense technology companies (e.g., growth-stage drone, missile, and maritime systems companies). Commercial cloud, AI, and imagery companies, for instance, can also provide high-quality solutions to allies and partners.

## Commercial Technologies and Modern Warfighting

Though the effects of drones tend to garner the most attention in the Russia-Ukraine war, Ukraine’s rapid adoption of commercial technologies highlights their importance to building resilient, adaptable militaries that are capable of holding the line against adversaries like China and Russia. Commercial cloud and edge computing proved essential for the Ukrainian government to withstand Russian cyberattacks at the outset of the war (Mitchell, 2022). Fortunately, Ukraine invested significantly in building national-level digital infrastructure—all built with commercial technologies—three years before Russia launched its invasion. The digital infrastructure was not for military or security applications per se. Rather, this infrastructure was designed to modernize civilian agency functions and the citizenry’s “user experience” vis-a-vis government services. For example, the digital infrastructure was used to modernize tax reporting, drivers licenses, and entitlement programs. President Zelensky framed this effort as building a “state in a smartphone” (Zelensky, 2019).

That very same infrastructure is now being used by an extraordinary ecosystem of companies, research institutions, and the Ukrainian government to rapidly develop new applications for defense, security, and intelligence functions. For example, apps were developed for Ukrainians to securely report Russian military positions and activities. Ukrainians are integrating commercial off-the-shelf software, data analytics, compute, and artificial intelligence with commercial products (e.g., sensors, small drones, and communications devices) to rapidly



innovate and push new capabilities into the field. Ukrainians are using those same commercial technologies to also optimize employment of munitions and other military systems. For example, commercially-available computer vision models and analytics are allowing Ukrainians to more precisely identify targets and optimize use of critical limited resources like 155-millimeter shells.

Ukraine is not alone in its adoption of commercial technologies to enhance military capabilities. For its part, the United States' military is attempting to do much the same. The fiscal year 2026 National Defense Authorization Act reinforced that the DoD should take a commercial first approach to acquisitions. President Trump's EO, "Modernizing Defense Acquisitions and Spurring Innovation in the Defense Industrial Base" states that the DoD will have a "first preference for commercial solutions" (White House, 2025).

China and Russia are also leveraging commercially available technologies at scale. China's military-civil fusion strategy is core to the People's Liberation Army's (PLA's) modernization efforts (Center for Security and Emerging Technology, 2025). As just one example, the PLA is integrating commercial artificial intelligence into virtually all of its operations—everything from backend enterprise systems, to logistics and command, control, communications, computers, cyber, intelligence, surveillance, and reconnaissance (C5ISR) (Center for Security and Emerging Technology, 2026). A February 2026 report from the Center for Security and International Studies (CSIS) noted how Russia is integrating open-source and commercial artificial intelligence models like Llama (produced by Facebook parent company, Meta) and Mistral (a French company) into its command and control (C2) architectures (Bondar, 2026).

## **Evolving Security Cooperation to Promote Commercial Solutions**

Helping allies and partners keep pace with technological change, build resilient forces capable of assuming greater ownership of security challenges, and enhancing collective deterrence against China and Russia requires the United States to offer a full suite of potential solutions. Non-program of record technologies are essential in that regard. The role that commercial solutions, specifically, can play in achieving these goals warrants more explicit attention.

To be clear, there are efforts already underway to work with allies and partners on commercial technology applications. For example, the Defense Innovation Unit (DIU) is administering prize challenges for AUKUS (DoD, 2024a). DoD teams are also testing commercial solutions for maritime security in Southeast Asia (DoD, 2024b). However, IDA's research found that the Do currently lacks a broader vision to scale and sustain such efforts (Lane, 2024). IDA's research also revealed four high-level insights that can help the Department close this gap.

First, defense innovation organizations (DIOs) like the DIU do not have the mandate, resources, or infrastructure to consistently support efforts to help allies or partners explore commercial solutions. Second, legal and institutional barriers within the Department inhibit the DIOs' ability to consistently share data with security cooperation entities like DSCA.

Third, as noted above, discernment will be required in terms of what commercial technologies and applications to prioritize. IDA's analysis suggested that commercial compute, data, cyber, and artificial intelligence (e.g., multi-modal and generative) can have the greatest transformative effect for allies and partners. Further, commercial technologies could be most impactful for allies and partners' ISR, measurement and signals intelligence, and maritime domain awareness capabilities, though there are many other high-impact applications worthy of consideration.



Finally, IDA posited multiple potential business models for how the Department could create the organizational structure needed to engage allies and partners on commercial technology adoption. One potential structure proposed by IDA was for DSCA to develop a dedicated team resourced and mandated to lead such efforts, drawing on technical expertise from across the Department.

Given the government's clear intent to open the aperture for allies and partners to acquire non-program of record solutions, this paper revisits and refreshes some of the recommendations first made by IDA in 2024. We will also offer new ideas for consideration as the security cooperation enterprise implements policy and organizational changes, all with the intent of helping position the enterprise to more effectively incorporate commercial technologies into security cooperation activities with allies and partners.

### **Ensure Commercial Technologies Factor into Organizational Mandates and Design**

Congress's recent mandate to create an office focused on non-program of record sales is generally consistent with one of IDA's recommendations to consider a dedicated team within DSCA focused on commercial technologies. Though it is unclear at this time where the office which Congress has mandated will ultimately reside within the Department, many of the practical steps that IDA offered for how DSCA could create such an office may still apply.

First, commercial technologies are an important and impactful subset of non-program of record systems. As such, commercial technologies should be a defined priority for the new non-program of record office. Second, to effectively promote and enable adoption of commercial technologies among allies and partners, the office will need to clearly define its value proposition to industry, allies, and partners alike.

The office could deliver value to commercial technologies and allies and partners by facilitating customer discovery and requirements generation. The office could build and establish relationships and deep understanding of allies and partner needs and help develop procurement strategies and roadmaps for the incorporation of commercial technologies into allied and partner operations. Further, the office could leverage its convening power and create network effects. Specifically, the office can serve as a point of entry for commercial companies interested in selling their technologies to foreign militaries. By leveraging its relationships and understanding of allied and partner requirements, this office can help new, commercial entrants to foreign military sales more effectively break into the market.

Finally, the office will need to cultivate and maintain a new ecosystem of partners, both in industry and government. Commercial technology companies and new entrants to the defense market often do not have the institutional knowledge or mechanisms needed to effectively engage foreign military customers. Helping them develop the internal capacity to do so will be important for the office to succeed in expanding the base of companies participating in the international arms sales market beyond the traditional primes. Consistent with President Trump's EO, "Establishing an America First Arms Transfer Strategy," the office should also work closely with the Department of State and the Department of Commerce to help commercial companies and new entrants navigate international defense markets and connect with customers.

### **Use Marketplaces of Vetted Commercial Technologies**

The DoD invests heavily in validating commercial technologies to ensure their trustworthiness, security, and reliability. It is also parlaying these validation efforts to create internally-facing marketplaces designed to accelerate procurement and adoption of commercial solutions across the Department and the Services. For example, the DIU created a Blue UAS List (which has since transitioned to the Defense Contract Management Agency). The list





includes drones that are compliant and approved by the Department for procurement. Platform One (led by the U.S. Air Force) effectively serves as a marketplace for DoD teams to obtain tools needed for software development. Opening such marketplaces to foreign customers (or creating a foreign-facing marketplaces) could help the security cooperation enterprise more quickly identify potential solutions and potentially accelerate solution delivery.

### **Modify Foreign Military Financing to Promote American Technology**

Currently, Foreign Military Financing (FMF) can only be used by Israel, Egypt, Jordan, Morocco, Tunisia, Turkey, Portugal, Pakistan, Yemen, and Greece to enter into a contract directly with a defense firm (i.e., a “direct commercial contract”, or DCC). Congress could amend legislation allowing any ally or partner to leverage FMF for a DCC, within some limits. For example, permission to use FMF for DCCs could be constrained to acquiring commercial solutions already included in one of the DoD’s aforementioned marketplaces.

Another potential constraint could be limiting use to the commercial technologies that the United States government wants to promote abroad the most (e.g., commercial computing service providers, artificial intelligence, or communications technologies) in the context of geopolitical competition with China and Russia.

### **Leverage Program Acquisition Executives’ Insights into Commercial Solutions**

Many Program Acquisition Executives (PAEs) are already exploring or incorporating commercial solutions into their programs. However, assuming legislation and policy are implemented as intended, PAEs will be increasing their level of engagement with commercial solutions providers. With DSCA’s realignment to USW(A&S), the security cooperation enterprise has an opportunity to glean insights into the commercial technologies being considered, piloted, and adopted across the Department’s acquisition portfolios. This can help the security cooperation enterprise identify vetted, trusted solutions for allies and partners’ operational needs.

### **Leverage Allies and Partners to Test Commercial Solutions**

As a counterpoint to leveraging PAEs to identify potential solutions, allies and partners may be able to help the Department’s acquisition community explore new technologies and applications, too. For example, by leveraging experimentation authorities provided in Section 873 of the fiscal year 2026 National Defense Authorization Act, Combatant Commands could sponsor experimentation and prototyping of commercial technologies in collaboration with allies and partners. The Combatant Commands and USMILGRPs can gather insights from allies and partners on the efficacy of such commercial technologies, helping inform the DoD’s own procurements.

### **Prepare the Security Cooperation Enterprise to Promote Commercial Solutions**

Combatant Commands, Service Component Commands, and USMILGRPs need to be informed on commercial technology applications and key considerations for how allies and partners can effectively adopt them. The Defense Security Cooperation Service and Defense Security Cooperation University can address this need by incorporating additional use cases and guidance into training and education, preparing members of the Defense Security Cooperation Service for considering and implementing different kinds of solutions to allies and partners defense challenges.

As the acquisition community adopts this new mission area, as well, they should be prepared for the unique requirements of allies and partners. This involves earlier integration of allies and partners potential requirements into planning, and considering exportability requirements earlier in the design and adoption of new technologies. This is a mandate for the



Warfighting Acquisition University, to open the aperture for acquisition professional to include potential sales outside of the United States in their certification and coursework.

## Final Thoughts

As the acquisition community assumes greater responsibility for oversight and execution of the Department's security cooperation mission, it is critical that security cooperation is not viewed exclusively through the lens of acquisitions and sales transactions. For example, assisting allies and partners in improving their institutional capacity (i.e., policy and strategy, resource management, human resource management, logistics management, military operations, joint concepts, and intelligence) will help ensure they are not acquiring capabilities for a hollow force.

Rather, if arms transfers and institutional capacity building efforts are conducted synergistically, the Department can best reinforce the U.S.'s status as preferred partner (thereby helping position the United States long-term as allies' and partners' key source of weapons and platforms) and enable other countries to take on greater burden sharing.

## References

- Bondar, K. (2026). *How Russia is reshaping command and control for AI-enabled warfare*. Center for Strategic and International Studies. <https://www.csis.org/analysis/how-russia-reshaping-command-and-control-ai-enabled-warfare>
- Center for Security and Emerging Technology. (2025). *Pulling back the curtain on China's military-civil fusion*. <https://cset.georgetown.edu/wp-content/uploads/CSET-Pulling-Back-the-Curtain-on-Chinas-Military-Civil-Fusion.pdf>
- Center for Security and Emerging Technology. (2026). *China's military AI wish list*. <https://cset.georgetown.edu/publication/chinas-military-ai-wish-list/>
- Defense Security Cooperation Agency. (2026). *Department of War finalizes realignment of the Defense Security Cooperation Agency*. <https://www.dsca.mil/Press-Media/Article-Display/Article/4402760/department-of-war-finalizes-realignment-of-the-defense-security-cooperation-age>
- DoD. (2023). *Department of Defense unveils comprehensive recommendations to strengthen foreign military sales*. <https://www.war.gov/News/Releases/Release/Article/3425963/department-of-defense-unveils-comprehensive-recommendations-to-strengthen-forei/>
- DoD. (2024a). *DIU partners with AUKUS Pillar II for international prize challenge*. <https://www.defense.gov/News/Releases/Release/Article/3718379/diu-partners-with-aukus-pillar-ii-for-international-prize-challenge/>
- DoD. (2024b). *DoD launches industry-government maritime security (MARSEC) consortium to advance maritime security across Southeast Asia*. <https://www.war.gov/News/Releases/Release/Article/3968018/dod-launches-industry-government-maritime-security-marsec-consortium-to-advance/>
- DoD. (2025). *Unifying the department's arms transfer and security cooperation enterprise to improve efficiency and enable burden sharing*. <https://media.defense.gov/2025/Nov/10/2003819440/-1/-1/1/UNIFYING-THE-DEPARTMENTS-ARMS-TRANSFER-AND-SECURITY-COOPERATION-ENTERPRISE-TO-IMPROVE-EFFICIENCY-AND-ENABLE-BURDEN-SHARING.PDF>



- DoD. (2026). *National defense strategy of the United States of America*.  
<https://media.defense.gov/2026/Jan/23/2003864773/-1/-1/0/2026-NATIONAL-DEFENSE-STRATEGY.PDF>
- Kurilla, M. E. (2025). *Posture statement of U.S. Central Command*.  
<https://www.centcom.mil/ABOUT-US/POSTURE-STATEMENT/>
- Lane, N. et al. (2024). *Building partner nation capabilities with commercial technology*. Institute for Defense Analyses. <https://www.ida.org/research-and-publications/publications/all/b/bu/building-partner-nation-capabilities-with-commercial-technology>
- Mitchell, R. (2022). How Amazon put Ukraine’s “government in a box”—and saved its economy from Russia. *Los Angeles Times*. <https://www.latimes.com/business/story/2022-12-15/amazon-ukraine-war-cloud-data>
- U.S. Congress. (2024). *Public Law 118–159*.  
<https://www.congress.gov/118/plaws/publ159/PLAW-118publ159.pdf>
- U.S. Department of State. (2025). *Fiscal Year 2024 U.S. arms transfers and defense trade*.  
<https://www.state.gov/fiscal-year-2024-u-s-arms-transfers-and-defense-trade/>
- U.S. House Committee on Foreign Affairs. (2023). *McCaul announces foreign military sales task force*. <https://foreignaffairs.house.gov/news/press-releases/mccaul-announces-foreign-military-sales-task-force>
- U.S. House Committee on Foreign Affairs. (2024). *Foreign military sales tiger task force report*.  
<https://foreignaffairs.house.gov/sites/evo-subsites/foreignaffairs.house.gov/files/migrated/uploads/2024/02/2.7.24-FOREIGN-MILITARY-SALES-TIGER-TASK-FORCE-REPORT.pdf>
- The White House. (2025a). *Modernizing defense acquisitions and spurring innovation in the defense industrial base*. <https://www.whitehouse.gov/presidential-actions/2025/04/modernizing-defense-acquisitions-and-spurring-innovation-in-the-defense-industrial-base/>
- The White House. (2025b). *National security strategy*. <https://www.whitehouse.gov/wp-content/uploads/2025/12/2025-National-Security-Strategy.pdf>
- The White House. (2025c). *Reforming foreign defense sales to improve speed and accountability*. <https://www.whitehouse.gov/presidential-actions/2025/04/reforming-foreign-defense-sales-to-improve-speed-and-accountability/>
- The White House. (2026). *Establishing an America First arms transfer strategy*.  
<https://www.whitehouse.gov/presidential-actions/2026/02/establishing-an-america-first-arms-transfer-strategy/>
- Zelenskyy, V. (2019). *I dream about a state in a smartphone*.  
<https://www.president.gov.ua/en/news/ya-mriyu-pro-derzhavu-u-smartfoni-volodimir-zelenskij-55585>



## PANEL 12. BEYOND THE LINEAR BARRIER: DATA-DRIVEN STRATEGIES FOR ACCELERATED ACQUISITION

Wednesday, May 6, 2026, 1615 – 1730 ET (1315 - 1430 PT)

### Panel Summary:

As threats and technology evolve at an unprecedented pace, the traditional linear acquisition model is being challenged by a global push for speed and iterative innovation. This panel provides a data-driven examination of recent progress in reducing cycle times, contrasting statistically significant improvements in Major Defense Acquisition Programs (MDAPs) with persistent structural inefficiencies identified by the Government Accountability Office (GAO). By analyzing leading product development practices and a proposed "Korean-style" rapid acquisition model, these researchers offer a roadmap for institutionalizing speed and ensuring that newer programs are structured for immediate operational impact.

**Chair: Major General Christopher D. Schneider, USA**, Deputy for Acquisition and Systems Management, ASA (ALT)

### Presenters:

**A Study on Establishing a Rapid Acquisition Process for the Military Application of Advanced Commercial Technologies in Korea** – *Won Joon Jang, Jeonbuk National University, South Korea*

**Recent Improvements in Acquisition Speed to Initial Operational Capability** – *Dr. Philip S. Anton, Stevens Institute of Technology*

**Weapon Systems Annual Assessment: DOD Leaders Should Ensure That Newer Programs Are Structured for Speed and Innovation** – *Erin Carson, U.S. Government Accountability Office (GAO)*



**Major General Christopher D. Schneider, USA**— received his commission as an Infantry Officer from the United States Military Academy in 1994. His tactical assignments include: Rifle and Reconnaissance Platoon Leader, Headquarters and Headquarters Company Executive Officer, and Battalion Assistant S-3, 1st Infantry Division; Assistant G-3 Training Officer for XVIII Airborne Corps; Airborne Rifle Company Commander and an Airborne Headquarters and Headquarters Company Commander in 3-325 A.I.R., 82d Airborne Division; and Ranger Assistant Operations Officer, 75th Ranger Regiment. He has multiple operational/combat deployments to Bosnia-Herzegovina, Iraq, and Afghanistan.

In 2005, Major General Schneider entered the Acquisition Corp and served in numerous assignments to include: (LTC/COL) Assignment Officer in the Human Resources Command, Alexandria, Virginia; Product Manager for Soldier Maneuver Sensors in PEO Soldier; Branch Chief for the Acquisition Branch in the Human Resources Command; Military Assistant to the Assistant Secretary of Defense for Acquisition; Project Manager for Soldier Sensors and Lasers and as the Assistant Program Executive Officer (APEO) for Futures and Integration in PEO Soldier. Project Manager for the Integrated Visual Augmentation System, a Rapid Prototyping initiative in PEO Soldier. He recently served as the Deputy for Acquisition and Systems Management, Office of the Assistant Secretary of the Army (Acquisition, Logistics and Technology).

Major General Schneider's military education includes the Infantry Officer Basic Course, Ranger School, Air Assault School, Bradley Leader Course, Scout Platoon Leader Course, Jumpmaster Course, Armor Officer Advanced Course, Combined Arms and Services Staff School, the Army Acquisition Basic Course, U.S. Army Command & General Staff College, and the National Defense University. Major General Schneider has received a B.S. in Mechanical Engineering from the United States Military Academy, a



M.B.A. from Liberty University, and a M.S. in National Resourcing and Strategy from the Eisenhower School.

His awards and decorations include the Defense Superior Service Medal, Legion of Merit, Bronze Star Medal, Meritorious Service Medal with four oak leaf clusters, Joint Commendation Medal, Army Commendation Medal with oak leaf cluster, Army Achievement Medal with two oak leaf clusters, National Defense Service Medal, Armed Forces Expeditionary Medal, Afghanistan Service Medal, Iraq Service Medal, Valorous Unit and Superior Unit Award, Senior Parachutist Badge, German Parachutist Badge, Ranger Tab, Air Assault Badge, Expert Infantryman Badge, and Combat Infantryman Badge.

Major General Schneider is a native of Louisville, Kentucky, and is married to COL Maria Schneider, an Army Acquisition Officer and native of Miami, Florida. They have three children: Michelle, Jacob, and Luke.





# **A Study on Establishing a Rapid Acquisition Process for the Military Application of Advanced Commercial Technologies in Korea**

**Won Joon Jang, PhD**—Associate Professor in the Department of Advanced Defense Technology and Industry at Jeonbuk National University (JBNU). He received his PhD in Economics from Seoul National University's Graduate School of Public Administration (Techno-Economics Policy Program). Prior to joining academia, he served as Director of the Defense Industry Research Division at the Korea Institute for Industrial Economics and Trade (KIET). He was also a Visiting Fellow at the Center for Strategic and International Studies (CSIS) in Washington, D.C. He currently serves as Vice President of the Korea Society of Innovation and holds advisory roles with government ministries and local governments. His research focuses on defense acquisition, defense industry strategy, and civil-military technology integration with more than 130 papers and reports. [wjjang47@jbnu.ac.kr]

## **Abstract**

This study proposes a new rapid acquisition procedure for the swift military application of advanced commercial technologies in Korea, reflecting the growing global trend of accelerating the integration of cutting-edge commercial innovations into defense systems. To this end, the research examines and analyzes the institutional frameworks and practical cases of rapid military adoption in leading countries—including the United States, Israel, and Ukraine—to identify key success factors, best practices, and limitations. In particular, a nationwide survey of experts from Korean industry, academia, research institutes, and the military was conducted to assess the shortcomings of Korea's traditional weapon acquisition processes and derive a rapid acquisition model tailored to the Korean context. Based on these findings, the study presents actionable Korean style rapid acquisition process, policy and institutional recommendations that can enhance Korea's ability to accelerate the defense adoption of advanced civilian technologies to a level comparable with leading global defense nations.

## **Introduction**

The contemporary security environment is characterized by unprecedented technological acceleration. Advances in artificial intelligence, unmanned systems, satellite communications, and big data analytics are reshaping the nature of warfare at a fundamental level. Unlike previous generations of military innovation, these technologies are largely developed in the commercial sector and subsequently adapted for military use.

Recent conflicts have demonstrated this transformation with striking clarity. In particular, the Russia–Ukraine War has shown that commercially available drones, civilian satellite networks, and open-source intelligence platforms can significantly influence battlefield outcomes. More recently, conflicts involving Iran and its regional engagements have further highlighted the increasing use of drones, missile systems, and networked surveillance capabilities, reinforcing the role of rapidly adaptable and often commercially derived technologies in modern warfare. These developments underscore the growing importance of speed in technology adoption and operational deployment.

Against this backdrop, traditional defense acquisition systems are increasingly inadequate. These systems were originally designed to support large-scale, capital-intensive weapon systems developed over long time horizons. However, such approaches are ill-suited for technologies that evolve rapidly and require continuous iteration.

South Korea faces a particularly acute challenge in this regard. While the country possesses advanced technological capabilities and a strong industrial base, its defense acquisition system remains largely rooted in conventional procurement models. Although recent



policy initiatives have sought to introduce elements of rapid acquisition, their impact has been limited due to deeper structural constraints.

This study addresses the following research questions. First, how can South Korea establish an effective rapid acquisition process capable of integrating advanced commercial technologies? Second, what institutional and procedural barriers hinder such integration? Third, what lessons can be drawn from leading countries, and how can these be adapted to the Korean context.

## **Research Framework and Methodology**

This study adopts a comprehensive research framework that integrates theoretical analysis, comparative case studies, and empirical investigation. The research is structured around three main components: analysis of global rapid acquisition models, evaluation of Korea's existing acquisition system, and development of a policy-oriented reform model.

First, the study conducts an in-depth review of global defense innovation trends, focusing on how advanced countries have restructured their acquisition systems to accommodate commercial technologies with velocity. Particular attention is given to the United States, Israel, and Ukraine, which represent distinct but complementary models of rapid acquisition.

Second, the study examines the structural limitations of Korea's acquisition system through institutional analysis. This includes an evaluation of procedural frameworks, governance structures, and incentive mechanisms.

Third, the study incorporates findings from a nationwide expert survey involving stakeholders from industry, academia, research institutes, and the military. This empirical component provides critical insights into the practical challenges and constraints faced by key actors within the defense ecosystem.

Finally, based on these analyses, the study proposes a Korean-style rapid acquisition model designed to address identified limitations while leveraging global best practices.

## **Global Trends in Rapid Defense Acquisition**

The increasing importance of commercial technologies in military applications is a defining feature of contemporary defense innovation. Unlike traditional defense technologies, which were primarily developed within government-controlled systems, many of today's most advanced technologies originate in the private sector.

The United States has been at the forefront of this transformation, particularly through the establishment of the Defense Innovation Unit (DIU) in the middle of 2010s. The DIU was created to bridge the gap between the Department of Defense and commercial technology companies, particularly those in Silicon Valley. Through mechanisms such as the Commercial Solutions Opening (CSO), the United States has developed a flexible acquisition process that enables rapid contracting and prototyping.

The CSO process typically involves identifying operational problems, soliciting solutions from commercial firms, rapidly awarding contracts, and conducting prototype development within a short timeframe. Unlike traditional procurement processes, which can take over a decade, CSO projects are often completed within one to two years. Empirical data indicate that a significant proportion of these projects successfully transition to operational deployment, demonstrating the effectiveness of this approach.



Israel offers another compelling model through its iHLS ecosystem, which emphasizes civil-military integration and startup-driven innovation. The iHLS framework supports the entire life cycle of innovation, from problem identification and startup selection to structured accelerator programs, prototyping, testing, and commercialization. In particular, the accelerator programs provide startups with mentoring, technical validation, access to military users, and connections to investors, thereby facilitating rapid capability development and market readiness. This model is characterized by its strong emphasis on real-world testing and rapid iteration, enabling technologies to be deployed quickly and effectively.

The Ukrainian experience, particularly during wartime, further highlights the importance of agility in acquisition systems. Faced with urgent operational needs, Ukraine has rapidly integrated commercial drone technologies and other innovations into its military operations. A central driver of this effort has been the establishment of the Brave1 platform, which serves as a core innovation and coordination hub connecting the government, military, and private-sector technology developers. Through this platform, Ukraine has been able to accelerate proposal evaluation, funding allocation, field testing, and deployment in an integrated manner. This demonstrates that under the right institutional conditions, rapid acquisition is not only possible but essential for operational effectiveness.

**Table 1. Comparative Overview of Rapid Acquisition Systems in Major Countries**

| Category              | USA(DIU)                                                                                 | Israel (iHLS)                                                      | Ukraine (Brave1)                                                          |
|-----------------------|------------------------------------------------------------------------------------------|--------------------------------------------------------------------|---------------------------------------------------------------------------|
| Launch Year           | 2015                                                                                     | 2012                                                               | 2023                                                                      |
| Governing Body        | U.S. Department of Defense (DoD)                                                         | Private sector + Ministry of Defense (SIBAT)                       | Ministry of Digital Transformation                                        |
| Key Technology Areas  | Commercial technologies (AI, robotics, cybersecurity, etc.)                              | Dual-use technologies (security, cyber, sensors)                   | Drones, cyber, and battlefield IT solutions                               |
| Operational Structure | Government-led with private-sector engagement                                            | Private accelerator-driven ecosystem                               | Wartime-oriented digital innovation platform                              |
| Contracting Mechanism | Flexible contracting using OTA (Other Transaction Authority)                             | PoC-based model linking demonstration to commercialization         | Fast-track approval with priority on battlefield deployment               |
| Core Characteristics  | Institutionalized rapid transition of commercial technologies into military capabilities | Private-sector-led defense innovation through accelerator programs | Civil-military integrated rapid innovation model under wartime conditions |

A comprehensive review of major countries' efforts to rapidly integrate advanced commercial technologies into military applications reveals a set of critical policy implications across multiple dimensions, including policy orientation, legal and institutional frameworks, organizational structures, budgeting mechanisms, acquisition processes, and supporting infrastructure.

From a policy perspective, leading countries exhibit distinct but converging strategic orientations toward rapid acquisition. The United States, through the DIU and its CSO mechanism, emphasizes the rapid military application of advanced commercial technologies by



prioritizing speed, flexibility, and direct engagement with non-traditional defense firms. In addition, the United States has institutionalized rapid acquisition through the use of Other Transaction Authority (OTA) and the establishment of the Adaptive Acquisition Framework (AAF), which includes pathways such as the Middle Tier of Acquisition (MTA)—enabling rapid prototyping and fielding typically within five years—and tailored Software Acquisition approaches for iterative and agile capability development. These mechanisms collectively demonstrate a successful model for integrating advanced commercial technologies into defense acquisition systems. In contrast, Israel's iHLS and related innovation platforms focus on expanding an open defense innovation ecosystem driven by private-sector creativity and startup participation. Ukraine's Brave1 initiative represents a wartime model, where the primary objective is the immediate deployment of operationally viable technologies, thereby accelerating innovation under battlefield conditions.

In terms of legal and institutional frameworks, flexibility emerges as a common denominator. The United States leverages OTA to bypass traditional Federal Acquisition Regulation (FAR) constraints, enabling agile contracting and rapid technology adoption. Israel operates under a relatively unconstrained legal environment, allowing private-sector-led innovation without significant regulatory barriers. Ukraine utilizes emergency procurement laws and digital platforms to expedite decision-making and resource allocation in wartime.

Organizationally, these countries have established dedicated entities to facilitate rapid acquisition. The United States has institutionalized innovation through the DIU and rapid acquisition offices within the military services, fostering collaboration among the military, industry, academia, and research institutions. Israel's model is characterized by private-sector leadership, with iHLS serving as a central hub connecting startups, investors, and defense stakeholders. Ukraine operates its Brave1 platform under the Ministry of Digital Transformation, highlighting the role of digital governance in defense innovation.

Budgetary arrangements further reinforce these institutional differences. The United States allocates substantial resources through its Research, Development, Test, and Evaluation (RDT&E) budget, with DIU funding complemented by additional defense appropriations for scaling successful projects. Israel combines government funding with private investment, providing targeted financial support to startups through defense R&D programs. Ukraine adopts a hybrid approach, mobilizing both public and private funding to support rapid innovation under resource-constrained conditions.

The acquisition processes themselves highlight the most significant divergence from traditional systems. The U.S. CSO model follows a streamlined sequence of needs identification, problem definition, solicitation, evaluation, contracting, prototyping, and transition to production, typically completed within one to two years. Israel's iHLS model operates on a life cycle approach encompassing needs identification, open calls, startup selection, acceleration, testing, investment linkage, and commercialization, often within six to twelve months. Ukraine's Brave1 platform enables an even faster cycle—from proposal submission to battlefield deployment—within two to six months.

Finally, infrastructure plays a critical enabling role in these systems. The United States maintains multiple DIU hubs across major innovation regions, supported by extensive human networks and venture capital linkages. Israel benefits from a concentrated technology hub in Tel Aviv, coupled with robust testing and validation facilities. Ukraine utilizes digital platforms and operational testbeds to accelerate deployment.

In summary, while each country adopts a distinct approach, several common principles emerge: flexibility in legal frameworks, strong institutional support, integration of private-sector innovation, rapid prototyping and testing, and alignment with national strategic objectives. These



findings provide important implications for designing a Korean-style rapid acquisition system capable of bridging the gap between commercial innovation and military capability.

**Table 2. Comparative Overview of Rapid Acquisition Systems in Major Countries**

| Category                               | United States (DIU)                                                                                                                                     | Israel (iHLS)                                                                                                                                 | Ukraine (Brave1)                                                                                                        |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Policy / Strategy</b>               | Rapid military application strategy leveraging commercial technologies                                                                                  | Expansion of an open defense innovation ecosystem driven by private-sector creativity                                                         | Rapid deployment of operationally viable technologies for battlefield use                                               |
| <b>Legal / Institutional Framework</b> | Use of OTA to streamline acquisition and bypass FAR regulations                                                                                         | Minimal legal constraints; privately driven, agreement-based operations                                                                       | Wartime emergency regulations with platform-based execution                                                             |
| <b>Organization</b>                    | DIU structure including MET, CET, and Portfolio PM; rapid acquisition units across military services (Army, Navy, Air Force, Marine Corps, Space Force) | Private-sector-led iHLS ecosystem with collaboration from Ministry of Defense and military                                                    | Platform-based organization under the Ministry of Digital Transformation                                                |
| <b>Budget</b>                          | Funded by DoD (approx. \$1 billion annually) with additional private-sector investment                                                                  | Primarily private investment with supplementary government funding                                                                            | Combined funding from government and allied support; multiple projects in the range of tens to hundreds of millions USD |
| <b>Acquisition Process</b>             | Needs identification → problem definition → solicitation → pitching → contracting → prototyping → transition to production                              | Needs definition → open call → startup selection → accelerator program & MVP development → validation → investment → commercialization/export | Proposal submission → evaluation → funding → pilot operation → battlefield deployment                                   |
| <b>Acquisition Timeline</b>            | Within 12–24 months                                                                                                                                     | Typically within 6 months                                                                                                                     | Typically within 2–6 months (accelerated under wartime conditions)                                                      |
| <b>Infrastructure</b>                  | Five major DIU hubs (e.g., Silicon Valley), supported by NSIN (network) and NSIC                                                                        | Technology hub centered in Tel Aviv with shared testing and validation facilities                                                             | Brave1 digital platform and Ministry of Defense testbeds                                                                |
| <b>Summary</b>                         | Benchmark model with well-established budget, organization, and procedures                                                                              | Private-sector-driven gateway for rapid technology integration supported by military collaboration                                            | Highly flexible, combat-oriented model enabling rapid deployment of AI, drones, and robotics                            |

## Current Status and Structural Limitations of Korea's Defense Acquisition System

In response to the growing need for the rapid integration of advanced technologies into military capabilities, South Korea has introduced several policy instruments aimed at





complementing its traditional, long-cycle defense acquisition system. Among these, the Rapid Pilot Project and the Rapid Requirements system represent the core institutional mechanisms designed to accelerate the adoption of emerging technologies. In parallel, the government has also implemented a range of defense industry support programs to facilitate the entry of private firms into the defense sector and to foster technological innovation.

The Rapid Pilot Project is primarily designed as a pre-acquisition experimentation mechanism that enables the development and testing of prototypes within a relatively short timeframe. Typically, the program operates within a five-year period, including approximately two years for prototype development and an additional two to three years for subsequent urgent requirement designation and potential transition to acquisition. However, compared to leading rapid acquisition mechanisms such as the CSO, the system does not incorporate a formal requirement determination at the initial stage. Instead, requirement decisions are made only after successful demonstration, which introduces uncertainty in the transition to full-scale acquisition.

The Rapid Requirements system, in contrast, is structured to support expedited acquisition based on formally identified urgent operational needs. Unlike the Rapid Pilot Project, it includes an explicit requirement determination process and follows a more structured testing and evaluation framework. This allows for relatively clearer pathways toward initial production and deployment. However, despite this formalization, the system remains constrained by elements of the traditional acquisition framework, particularly in later stages of mass production and long-term force integration. Moreover, since its introduction in 2023, there have been no confirmed cases of actual program implementation, raising concerns regarding its practical effectiveness and operational viability.

1

Both systems share certain common characteristics, including a maximum project duration of approximately five years and the possibility of producing minimum operational quantities. However, significant differences exist in terms of requirement determination, evaluation processes, and pathways to full-scale production. In particular, the transition from prototype development to operational deployment remains ambiguous in both systems, especially with regard to follow-on production beyond initial batches. In many cases, subsequent large-scale acquisition still relies on conventional procurement processes, thereby limiting the effectiveness of rapid acquisition mechanisms.

---

<sup>1</sup> In response to these limitations, the Defense Acquisition Program Administration (DAPA) has recently initiated efforts, as part of a national policy agenda, to revise the proposed Advanced Defense Capability Acquisition Act (tentative title). These efforts aim to establish a new acquisition pathway for the rapid fielding of advanced commercial technologies within the next three years, reflecting a growing institutional recognition of the need for more agile and flexible acquisition mechanisms.



**Table 3. Comparison of Korea's Rapid Acquisition Systems**  
(Adapted from Jang et al., 2024)

| Category                                            | Rapid Pilot Project                                                              | Rapid Requirements                                                                  | Remarks                                                                                                                                                           |
|-----------------------------------------------------|----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Requirement Determination                           | X*                                                                               | O                                                                                   | *Urgent requirement can be determined after successful pilot project completion                                                                                   |
| Testing & Evaluation                                | △<br>Military pilot operation + performance verification testing)                | O<br>(Formal military testing and evaluation)                                       | Since 2024, performance verification testing may substitute formal T&E for systems developed through pilot projects (Defense Acquisition Program Act, Sept. 2023) |
| Minimum Operational Unit Production                 | O                                                                                | O                                                                                   | Minimum operational quantity can be produced through sole-source contracts with prototype developers (Defense Acquisition Act, Sept. 2023)                        |
| Initial Deployment / Follow-on Production (Phase 1) | Unclear                                                                          | O                                                                                   | Possible if initial production quantities are reflected during urgent requirement determination                                                                   |
| Full Deployment / Follow-on Production (Phase 2+)   | Unclear                                                                          | Unclear                                                                             | Conventional acquisition process (PPBEES) may apply (requirement determination → mid-term planning → feasibility study → production)                              |
| Project Duration                                    | Within 5 years (2 years for pilot + 2–3 years after urgent requirement decision) | Within 5 years                                                                      | —                                                                                                                                                                 |
| Budget                                              | Approx. KRW 58 billion (2024); typically tens of billions KRW per project        | No fixed budget limit (projects exceeding KRW 50 billion require feasibility study) | —                                                                                                                                                                 |

Despite these institutional efforts, Korea's rapid acquisition system continues to exhibit significant structural limitations when compared to leading international models. One of the most critical challenges is the lack of flexibility in contracting and evaluation procedures. Unlike systems such as the U.S. OTA, which enable highly adaptive contracting mechanisms, Korea's acquisition framework remains bound by rigid regulatory structures, limiting the ability to rapidly engage with commercial technology providers.

A second major limitation lies in the absence of a fully developed demonstration-based procurement system. While the Rapid Pilot Project includes elements of prototype testing, it does not provide a clear and consistent pathway for transitioning successful prototypes into operational deployment. This gap between experimentation and acquisition significantly reduces the overall effectiveness of the system.

Third, the linkage between acquisition programs and operational military users remains insufficient. Effective rapid acquisition systems require close and continuous interaction between developers and end-users, enabling iterative feedback and real-time adaptation. However, in Korea's current system, such integration is limited, resulting in a disconnect between technological development and operational requirements.

Fourth, barriers to entry for private-sector participants—particularly startups and non-traditional defense firms—remain high. Complex regulatory requirements, limited incentives, and uncertainties regarding intellectual property rights discourage participation from innovative commercial actors. As a result, the defense innovation ecosystem remains relatively closed compared to those of leading countries.

These structural limitations highlight the need for both incremental improvements and more fundamental institutional reform. In the short term, continued refinement of existing mechanisms—such as regulatory adjustments, expanded budgets, and improved operational guidelines—can help mitigate some of the current constraints. However, such measures alone are unlikely to achieve the level of performance observed in advanced countries.

Accordingly, there is a growing consensus on the need to establish a new acquisition pathway specifically designed for the rapid military application of advanced commercial technologies. One potential approach is the introduction of a Korean-style rapid acquisition mechanism, tentatively referred to as a “K-DIU CSO,” modeled on the U.S. DIU framework. Such a system would incorporate flexible contracting, problem-driven acquisition, and rapid prototyping cycles.

In addition, it is necessary to develop dedicated programs that facilitate the transition of research and development outcomes into fieldable prototypes. This could take the form of a Rapid Force Integration Program that bridges the gap between technology development and operational deployment. Furthermore, existing defense industry support programs should be more closely integrated with rapid acquisition processes to ensure that their outputs contribute directly to military capabilities.

In conclusion, while South Korea has made important strides in introducing rapid acquisition mechanisms, the current system remains constrained by structural and institutional limitations. Addressing these challenges will require not only policy adjustments but also the establishment of a fundamentally new acquisition framework capable of supporting rapid, flexible, and innovation-driven defense development.

## **Empirical Findings: Key Results from the Survey on Rapid Acquisition**

The survey targeted a diverse group of respondents, including experts from academia, industry, research institutes, and the military. In addition, companies selected under the “Defense Innovation 100” program were included to capture perspectives from emerging defense-related firms. A total of 58 valid responses were collected, comprising 16 expert participants and 42 representatives from industry from July 31 to August 29, 2025.

## **Perceived Limitations of the Traditional Acquisition System**

The survey results reveal a strong consensus that South Korea's traditional defense acquisition system—commonly referred to as the PPBEES (Planning, Programming, Budgeting,



Execution, Evaluation, and Sustainment) framework—is not adequately equipped to respond to the rapid evolution of advanced commercial technologies.

Approximately 70.2% of respondents evaluated the system's responsiveness as insufficient or very insufficient, while only 5.3% considered it effective. When converted to a five-point Likert scale, the average score was 2.26, indicating a below-average level of performance. These findings suggest that the current acquisition system is structurally misaligned with the pace of technological change, particularly in areas driven by commercial innovation.

### **Strong Demand for Acquisition System Reform**

In contrast to the negative evaluation of the existing system, there is overwhelming agreement on the necessity of reform. A total of 94.4% of respondents indicated that improvements to acquisition procedures are necessary, with 59.3% stating that such reforms are “very necessary.” The average score of 4.50 further underscores the urgency of institutional change. This perception is particularly strong among experts, whose average score reached 4.81, compared to 4.37 among industry respondents. This gap suggests that those with deeper policy and institutional experience perceive a greater need for systemic transformation.

### **Key Reasons for Reform: Structural Misalignment with Technological Change**

The survey identifies several key reasons why reform is considered necessary. The most frequently cited reason, accounting for 41.6% of responses, is that the traditional acquisition process is fundamentally incapable of supporting the rapid military application of advanced commercial technologies. Additional factors include the inability of the current system to keep pace with fast technological advancements (26.6%), structural barriers to participation by startups and small firms (14.5%), and the lack of institutional pathways for integrating commercial innovations into military systems.

### **Comparative Assessment with Advanced Countries**

When benchmarked against leading countries, Korea's rapid acquisition systems—specifically the Rapid Pilot Project and Rapid Requirements—are perceived as significantly underdeveloped. For the Rapid Pilot Project, 87.5% of respondents rated Korea's system as inferior to those of advanced countries, while only 12.3% considered it comparable. The results are even more pronounced for the Rapid Requirements system, with 91.1% indicating that it falls short of global standards. These findings clearly indicate that Korea's current approach is not competitive in the global context and requires substantial improvement.

### **Structural Weaknesses in Existing Rapid Acquisition Programs**

The survey further identifies specific areas requiring improvement within Korea's rapid acquisition programs. For the Rapid Pilot Project, the most critical issue is the lack of guaranteed transition to mass production following successful prototype development (20.5%). This is followed by insufficient incorporation of military requirements in early stages (14.6%), reliance on traditional defense contractors (12.3%), and the absence of effective transition mechanisms to full-scale acquisition (11.1%).

Similarly, for the Rapid Requirements system, the absence of dedicated budgets is identified as the most significant limitation (20.9%), followed by weak linkage to follow-on production (17.4%), limited profitability due to constrained production volumes (15.1%), and continued dominance of established defense firms (13.4%). These results highlight a common structural issue: the lack of a seamless pathway from innovation to deployment.

### **Major Barriers to Rapid Military Adoption**

The survey also identifies key barriers to the rapid military application of advanced commercial technologies. The most significant concern is insufficient budget allocation, which



received the highest score (4.20), followed closely by the lack of guaranteed military procurement after successful project completion (4.12). Other major barriers include inadequate legal frameworks (e.g., absence of OTA-like mechanisms), limited collaboration between developers and military users, insufficient reflection of operational needs, and restrictive security regulations. These findings suggest that financial uncertainty and institutional constraints are the primary obstacles to effective rapid acquisition.

### **Policy Priorities for Reform**

Finally, the survey identifies key priorities for improving acquisition procedures. The most important priority is the establishment of dedicated budgets for rapid acquisition programs, with an average score of 4.40. This is followed by guaranteed procurement of successful prototypes by specific military units (4.22), the introduction of demonstration-based evaluation systems (4.14), and mechanisms for rapid performance upgrades after initial deployment (4.06). Additional priorities include strengthening collaboration between developers and military users, aligning technological development with operational needs, establishing dedicated legal frameworks, and improving organizational capacity.

### **Summary of Findings**

Overall, the survey results reveal a consistent pattern. First, the existing acquisition system is widely perceived as inadequate. Second, there is strong consensus on the need for reform. Third, the key challenges are structural rather than technological. Finally, effective reform requires a comprehensive approach that addresses legal, institutional, financial, and operational dimensions. These findings provide a strong empirical foundation for the development of a Korean-style rapid acquisition model and underscore the urgency of institutional transformation.

### **Establishing a Korean-Style Rapid Acquisition Model**

To address the structural limitations identified in Korea's current defense acquisition system, this study proposes the establishment of a dedicated rapid acquisition pathway designed specifically for the military application of advanced commercial technologies. This proposed framework represents a fundamental departure from existing mechanisms such as the Rapid Pilot Project and Rapid Requirements system.

The Korean-style Rapid Acquisition model is designed to overcome the inherent limitations of traditional acquisition processes by introducing a flexible, problem-driven, and innovation-oriented approach. Drawing on best practices from leading international models—including the U.S. DIU, Israel's iHLS ecosystem, and Ukraine's Brave1 initiative—the proposed framework is tailored to Korea's institutional and industrial context.

At its core, the Korean style Rapid Acquisition model consists of a four-stage acquisition process that integrates rapid problem identification, flexible contracting, iterative prototyping, and scalable deployment.

#### **Stage 1: Military Needs Identification and Open Solicitation (Within 30 Days)**

The first stage focuses on identifying operational needs and translating them into actionable problem statements. Unlike traditional requirement-driven approaches, this stage emphasizes a problem-based framework in which military units articulate operational challenges rather than predefined technical specifications.

Once needs are identified, they undergo a process of validation and refinement through close coordination between operational units and acquisition authorities. This includes redefining problem statements, conducting market research, evaluating technological feasibility, and assessing potential risks.





Based on this analysis, a solicitation for solutions is issued to the private sector. This solicitation is designed to attract a broad range of participants, including startups, small and medium-sized enterprises, and non-traditional defense contractors. The entire process is structured to be completed within 30 days, ensuring rapid engagement with the commercial technology ecosystem.

### **Stage 2: Company Selection and Contracting (Within 90 Days)**

The second stage involves the evaluation of proposed solutions and the selection of participating companies. This process is designed to be significantly more streamlined and flexible than conventional procurement procedures.

Initially, companies submit concise solution proposals within a short timeframe, typically within 10 days. These proposals are evaluated based on criteria such as technological maturity, market readiness, security considerations, and implementation feasibility.

Shortlisted candidates are then invited to present their solutions through structured presentations or pitch sessions. The evaluation process involves a multidisciplinary team, including program managers, technical experts, military end-users, and external specialists.

Following this process, selected companies are invited to submit refined proposals, after which final selection decisions are made. Contracts or cooperative agreements are then concluded within 90 days, significantly reducing the time required for initial engagement.

### **Stage 3: Prototype Development and Operational Experimentation (Within Two Years)**

The third stage focuses on the development of prototypes and their validation under realistic operational conditions. This stage is critical for bridging the gap between conceptual innovation and practical military application.

Selected companies are provided with comprehensive support, including access to advisory panels or accelerator programs. These support mechanisms may include technical mentoring, proof-of-concept (PoC) development assistance, investor linkage, and regulatory guidance.

A key feature of this stage is the integration of military users into the development process. Prototype developers are matched with designated test units, and dedicated testbeds are established to facilitate continuous feedback and iterative improvement. In the case of AI-based systems, data sharing, model training, and operational feedback are actively supported to enhance performance.

In addition, mechanisms such as “Live Demo Days” are introduced to enable direct interaction between developers, military personnel, public agencies, and potential investors. These events serve as platforms for demonstrating capabilities, gathering feedback, and facilitating future scaling.

Operational experimentation is conducted through field testing in military environments, allowing for the evaluation of usability, effectiveness, and deployment readiness. Based on these results, decisions are made regarding the suitability of the technology for operational use.

### **Stage 4: Initial Deployment, Rapid Iteration, and Scaling (Within Two to Five Years)**

The final stage focuses on transitioning successful prototypes into operational deployment and expanding their application.

Initial deployment is conducted at the unit level, typically within the military unit that first identified the operational need. A minimum quantity of the system is procured to enable real-world usage and validation.



Following deployment, continuous feedback from operational users is used to drive rapid iteration and performance improvement. Existing funding mechanisms, such as those supporting the enhancement of current force capabilities, can be leveraged to support these activities.

Scaling is achieved through multiple pathways. First, formal requirement generation can be initiated to support full-scale production under existing acquisition frameworks. Second, the application of the technology can be expanded to other public agencies, such as police, coast guard, and emergency services, thereby increasing demand and utilization.

Third, partnerships with major defense contractors can be established to facilitate integration into larger systems and enhance production capacity. Fourth, private investment can be mobilized through venture capital and defense-focused funding mechanisms to support commercialization and growth.

Finally, export opportunities are actively pursued through participation in international defense exhibitions, government-to-government (G2G) cooperation programs, and engagement with key partner countries. Participation in multinational exercises and joint demonstrations further enhances the visibility and credibility of developed technologies.

**Table 4. Korean-Style Rapid Acquisition Process (Draft)**

| Phase                                                       | Sub-Stage                         | Key Activities                                                                                                         | Responsible Entity                                       | Timeline          | Number of participating companies |
|-------------------------------------------------------------|-----------------------------------|------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|-------------------|-----------------------------------|
| <b>Phase 1</b><br>Needs Identification & Problem Definition | 1-1 Needs Identification          | Identification and submission of urgent operational needs from military units                                          | Military units, K-DIU                                    | Within 30 days    | 5-100+                            |
|                                                             | 1-2 Problem Definition            | Refinement of military needs into problem statements aligned with advanced technologies                                | K-DIU PM, military, technical team                       |                   |                                   |
|                                                             | 1-3 Market Research & Feasibility | Market analysis, TRL assessment, feasibility and risk analysis                                                         | K-DIU PM, technical team, external experts               |                   |                                   |
|                                                             | 1-4 Solicitation                  | Development and announcement of solution solicitation (Request for Proposal)                                           | K-DIU PM                                                 |                   |                                   |
| <b>Phase 2</b><br>Company Selection & Contracting           | 2-1 Proposal Submission           | Submission of short proposals (≤ pages or ≤ slides)                                                                    | Participating firms                                      | Within 60-90 days | 5-20                              |
|                                                             | 2-2 Evaluation & Pitch            | Technical evaluation and presentation (multi-stage review)                                                             | K-DIU PM, technical team, military users, experts        |                   |                                   |
|                                                             | 2-3 Final Selection & Contract    | Final selection and contract/agreement signing                                                                         | K-DIU PM, contracting authority                          |                   |                                   |
| <b>Phase 3</b><br>Prototype Development & Experimentation   | 3-1 Prototype Development         | PoC support, test unit matching, Live Demo Day, investor linkage                                                       | Selected firms, military test units, advisors, investors | Within 1-2 years  | 1-5                               |
|                                                             | 3-2 Field Experimentation         | Operational testing and validation in military environments                                                            | Military users, K-DIU                                    |                   |                                   |
| <b>Phase 4</b><br>Deployment, Iteration & Scaling           | 4-1 Initial Deployment            | Initial procurement and deployment at requesting military unit                                                         | Military units                                           | Within 2-5 years  | 1-2                               |
|                                                             | 4-2 Rapid Iteration               | Performance improvement based on user feedback and operational use                                                     | Military, K-DIU, developers                              |                   |                                   |
|                                                             | 4-3 Scaling & Expansion           | Transition to full-scale production (PPBEES), expansion to public sector, private partnerships, investment, and export | DAPA, military, K-DIU, industry, investors               |                   |                                   |

## Conclusion and Policy Implications

This study has examined the structural limitations of South Korea's defense acquisition system in the context of rapidly evolving advanced commercial technologies and proposed a comprehensive framework for institutional reform. Drawing upon comparative analysis of leading countries and empirical findings from a nationwide survey, the study identifies a clear and urgent need for a paradigm shift in acquisition policy and practice.

The analysis demonstrates that the primary challenge facing Korea is not a lack of technological capability, but rather the inadequacy of its institutional framework. Existing mechanisms, including the Rapid Pilot Project and Rapid Requirements system, represent important initial efforts but remain constrained by structural limitations in flexibility, scalability, and integration. As a result, they are insufficient to meet the demands of rapid technological change and evolving operational requirements.

In response to these challenges, this study proposes the establishment of a Korean style rapid acquisition model. This framework represents a fundamental departure from traditional acquisition models, emphasizing a problem-driven, flexible, and innovation-oriented approach. The proposed four-stage process—comprising rapid needs identification, streamlined contracting, iterative prototyping, and scalable deployment—provides a structured yet adaptable pathway for integrating advanced commercial technologies into military capabilities.

However, the successful implementation of this framework requires more than procedural reform. It necessitates a comprehensive transformation across multiple dimensions, including policy, legal frameworks, organizational structures, budgeting mechanisms, and supporting infrastructure.

At the policy and strategic level, rapid acquisition must be elevated to a central pillar of Korea's advanced defense innovation strategy. This includes the articulation of a clear national vision, the prioritization of key technological domains, and the development of a long-term roadmap supported by pilot programs and phased expansion. The establishment of a dedicated organization, tentatively referred to as K-DIU, should serve as the institutional anchor for this transformation.

From a procedural perspective, the introduction of the Korean style Rapid Acquisition Model is essential. By adopting a four-stage acquisition process with clearly defined timelines, Korea can significantly reduce acquisition cycles while maintaining operational effectiveness. This approach also enables greater participation from non-traditional actors, thereby expanding the defense innovation ecosystem.

Legal and institutional reform is equally critical. New legislative measures should be introduced to support rapid acquisition, including the establishment of flexible contracting mechanisms, simplified approval processes, and demonstration-based evaluation systems. Existing laws, such as the Defense Acquisition Program Act and the Defense Industry Development Act, should be revised to accommodate these changes. In addition, consideration should be given to enacting a comprehensive legal framework for advanced defense strategic industries.

Organizational reform is required to ensure effective implementation. This includes the creation of dedicated rapid acquisition units within both central acquisition authorities and operational commands, as well as the expansion of existing organizations responsible for defense industry support. Specialized personnel and program management capabilities must be developed to support the new system.

From a financial perspective, the establishment of dedicated budget mechanisms is essential. A separate budget line for rapid acquisition should be introduced to ensure stable



funding, complemented by mechanisms to attract private investment, such as a defense-focused venture fund.

Infrastructure development is another key enabler. This includes the establishment of dedicated testbeds for operational experimentation, the development of digital platforms to facilitate rapid acquisition processes, and the creation of communication hubs to enhance collaboration among stakeholders.

Furthermore, it is essential to strengthen the linkage between defense industry support programs and acquisition systems. The introduction of a Rapid Force Integration Program can facilitate the transition of R&D outcomes into fieldable prototypes, while the integration of rapid acquisition procedures into existing support programs can enhance their practical impact.

In conclusion, the transition toward a rapid, flexible, and innovation-driven acquisition system is not merely an option but a strategic necessity for South Korea. As the pace of technological change continues to accelerate, the ability to rapidly integrate advanced commercial technologies into military capabilities will become a decisive factor in national security and defense competitiveness. The Korean-style Rapid Acquisition model provides a viable and actionable pathway for achieving this transformation, but its success will depend on sustained commitment to institutional reform and strategic execution.

## Acknowledgements

This paper is based on a research project titled “*A Study on Improving Acquisition Procedures for the Rapid Military Application of Advanced Commercial Technologies*,” which was conducted in 2025 as a commissioned research project by the Korea Research Institute for Defense Technology Planning and Advancement (KRIT). The contents of this paper have been revised and further developed based on the findings of that study. This paper is prepared as a working paper for presentation purposes and is not intended for journal publication. The views, findings, and conclusions expressed in this paper are solely those of the author and do not necessarily reflect the official positions of KRIT or any affiliated institutions.

## References

- Brave1. (2025). <https://brave1.gov.ua/en/>
- Defense Acquisition Program Administration. (2023a, May 9). *Accelerating weapon system acquisition: Reducing procurement timelines by half*. Press release.
- Defense Acquisition Program Administration. (2024). *Introduction of acquisition fast track and future directions*.
- Defense Acquisition Program Administration. (2023b). *Defense industry development master plan (2023–2027)*.
- Defense Industry Promotion Agency (Korea Institute for Defense Analyses affiliated body). (2024). *Overview of defense industry promotion programs and their application to acquisition projects*.
- Defense Innovation Unit. (2021, May 5). *How to pitch the DIU* [Video]. YouTube. <https://www.youtube.com/watch?v=0ivj2LGPMYk&t=26s>
- Defense Innovation Unit. (2024). *DIU FY2023 annual report*.
- Defense Innovation Unit. (2025). *Open solicitations*. <https://www.diu.mil/work-with-us/open-solicitations>



Defense Rapid Acquisition Center. (2024). *Trends and future challenges of rapid pilot projects*.

iHLS Accelerator. (2025). <https://accelerator.i-hls.com/>

Jang, W. J., et al. (2024). *Establishing a Korean-style defense acquisition framework in response to global trends*. Korea Institute for Industrial Economics and Trade (KIET).

Ministry of National Defense & Defense Acquisition Program Administration. (2024). *Defense science and technology innovation promotion act*.

Ministry of National Defense & Defense Acquisition Program Administration. (2025a). *Defense acquisition program act*.

Ministry of National Defense & Defense Acquisition Program Administration. (2025b). *Enforcement decree of the defense acquisition program act*.

Ministry of National Defense & Defense Acquisition Program Administration. (2025c). *Defense industry development and support act*.

PitchBook. (2025). <https://pitchbook.com/profiles/investor/265045-69>

Presidential Commission on State Affairs Planning. (2025). *Five-year national governance plan of the Lee Jae-myung administration*.

The Times of Israel. (2019, August 12). *Come test your tech with us, security establishment urges innovative startups*.





# Recent Improvements in Acquisition Speed to Initial Operational Capability

**Philip S. Antón, PhD**—is the Chief Scientist of the Acquisition Innovation Research Center (AIRC) in the Systems Engineering Research Center (SERC)—a multi-university UARC led by the Stevens Institute of Technology. Dr. Antón assesses defense needs and envisions and develops innovative applied research, prototyping, and piloting to facilitate application to acquisition practices, policies, and workforce development. Prior to AIRC, Dr. Antón was a Senior Information Scientist at the RAND Corporation for 23 years. From 2011–2016, he directed Acquisition Policy Analysis for USD(AT&L). His PhD and MS are in information and computer science from the University of California, Irvine, specializing in computational neuroscience and artificial intelligence. His BS in engineering is from UCLA. [panton@stevens.edu]

## Abstract

Threats and technology are changing at a faster pace, so defense acquisition needs to respond quickly to meet warfighter and operator needs. With emphasis on speeding acquisition, has delivery of initial capabilities gotten faster? Analysis of the time between formal initiation of Major Defense Acquisition Programs (MDAPs) at Milestone (MS) B or C to achieved Initial Operating Capability (IOC) found statistically significant reductions since 2010. Using Selected Acquisition Report (SAR) and Modernized SAR (MSAR) schedule data, analysis found that MDAPs initiated at MS B since January 2010 took an average of 5.6 years compared to 8.7 years before 2010. MDAPs that initiated at MS C (i.e., using relatively mature technology and thus skipping MS B—a best practice for speed) reached IOC in an average of 3.5 years overall with no statistical difference before or since 2010. Combining these cases, programs started at either MS B or C since January 2010 took an average of 5.1 years compared to 8.1 years before 2010. While some caution is warranted given the relatively low population sizes, these results indicate that the Warfighting Acquisition System has made statistically significant progress in speeding initial capabilities to the warfighter.

## Introduction

With the advent of increasing technological advancements and faster pace of increasingly capable threats, the United States has increasingly emphasized speed of acquiring and delivering capabilities to warfighters and operators.<sup>1</sup>

Perhaps the most cited speed measure is the time it takes for Major Defense Acquisition Programs (MDAPs) to field an Initial Operating Capability (IOC). Unfortunately, the most common view is to cite MDAP delivery time considering the population of all available MDAPs going back decades. At best, that may be the view shown in *Figure 1* with averages and median at about 7.6 years, although there are tendencies to cite the worst extremes and ignore those cases faster than the central tendency. For example, while historically the worst case was a statistical outlier at 21.2 years (16.3 years would be the maximum if the outlier was removed), there are examples on the other end of the spectrum that reached IOC in under a year. Also, one-fourth of MDAPs reach IOC in under 5 years (the target for middle-tier programs), half under 7.5 years, and three-quarters under 10 years.

While helpful, the broad view in *Figure 1* gives no insight into whether defense acquisition (especially at the MDAP level, where data are more widely available) is improving in speed. The large number of old MDAPs and the low rate at which new MDAPs are started makes it rather impossible to measure any trends to improve speed by just looking at the complete population as a whole.

---

<sup>1</sup> See, for example, Executive Order No. 14265, 2025; Secretary of Defense, 2017, 2022; Secretary of War, 2025a, 2025b, 2026; Under Secretary of Defense for Acquisition, Technology, and Logistics, 2014.



In response, this paper summarizes statistical analysis to compare the population of older MDAPs to those started more recently to see if the speed of acquiring MDAPs has changed given the more recent emphasis on speed.




**Figure 1. Time from MDAP Initiation at MS B or C to Achieved IOC**

## Data Sources and Analytical Approach

Here we use schedule data for MDAPs as recorded in SARs and MSARs available in the War Department's Advana database as of April 2025. This includes 128 MDAPs that have achieved MS B or C as early as June 1979 and as late as May 2023. Additional MDAPs that have not achieved IOC yet were excluded. The MDAP portfolio categories used below are those created by the department in 2025 for internal uses. Officially, MDAPs are formally established with their cost and schedule baselined at Milestone (MS) B,<sup>2</sup> but some programs skip MS B and are initiated at MS C. This is also the point when SARs (and more recently MSARs) are provided to Congress. Note that within and across capability areas, the dollar size and complexity of MDAPs can vary widely, but we examine all MDAPs as a group to identify any trends.

MDAPs are relatively few in number and started relatively infrequently compared to smaller acquisition programs. Therefore, to identify any recent improvements, we selected January 2010 as a dividing time between older and more recent MDAPs with the hopes of obtaining sufficient population sizes to find any statistical differences. This selection was somewhat arbitrary but followed the passage of the Weapon Systems Acquisition Reform Act of

<sup>2</sup> See the "Program Initiation" entry in the Defense Acquisition University (DAU) Glossary (DAU, 2026).

2009 that renewed emphasis on rigor in managing acquisition programs. Other dates could be chosen, but risk having insufficient population sizes for statistical tests.

In general, IOC is attained “when some units and/or organizations in the force structure scheduled to receive a system have received it and have the ability to employ and maintain it” (DAU, 2026). The period between initiation and IOC generally includes system development, test, and initial production to deliver an initial operational (fielded) capability. In some cases, we used program events that were analogous to IOCs. Also, in the few cases where an MDAP schedule had multiple IOCs, we used the earliest achieved IOC to measure the speed to delivering some initial capability.

## Statistical Tests

To compare cycle times between two populations of MDAPs before and since January 2010, we used the Mann Whitney test (also known as Wilcoxon rank sum test with continuity correction). Comparisons with p-values above 10% were deemed not statistically significant. We noted any instances with very weak statistical results (p-values between 5%–10%) but only deemed statistical significance for p-values <5%. Any outliers were identified using the Tukey algorithm to identify any cases that were above the third quartile (Q3) by over 1.5 times the inner-quartile range (i.e., greater than  $Q3 + 1.5*(Q3-Q1)$ ).

## Graphics

Figures below showing time from MDAP initiation to IOC are box-and-whisker charts, with the minimum (min.) at the bottom of the vertical line, a box from Q1–Q3 and a line for the median, and “X” at the average (avg.: the arithmetic mean), and the maximum (max.) at the top of the vertical line. Any outliers are identified with small circles. Numbers on the charts are averages. Any statistically significant differences between the population of MDAPs initiated before 2010 and since January 2010 are flagged with an asterisk (\*). Numeric values are provided in tables below the charts.


## Results

Let us first examine the population of MDAPs that were initiated at MS B (i.e., leaving out those that skipped MS B). We will then separately examine the MDAPs that started at MS C (i.e., skipped MS B), followed by comparing the full set of MDAPs that started at either MS B or C.

### MDAPs Initiated at MS B

*Figure 2* shows the time to actual (achieved) IOC for MS B–initiated MDAPs. The results show statistically significant improvements (reductions) in cycle time for all MDAPs overall as well as for MDAPs in the Aircraft/Aviation and Ship/Maritime capability areas for MS B–initiated MDAPs initiated at MS B since January 2010 compared to those initiated before calendar year (CY) 2010. Overall, average cycle time dropped from 8.7 years to 5.6 years. Average Aircraft/Aviation times dropped from 10.4 years to 7.0 years, and average Ships/Maritime times dropped from 8.8 years to 4.3 years. Medians were similar (see the table below the chart). Any apparent differences for the other capability areas are not statistically significant—in many cases because of very small population sizes.





| Capability Area   |              | Min. | Q1  | Median | Avg.  | Q3   | Max. | IQR | St.Dev. | N = |
|-------------------|--------------|------|-----|--------|-------|------|------|-----|---------|-----|
| Overall           | Before 2010  | 1.1  | 6.3 | 8.4    | 8.7   | 10.9 | 21.2 | 4.6 | 3.6     | 93  |
|                   | Since 2010 * | 0.8  | 3.7 | 5.3 *  | 5.6 * | 7.7  | 9.5  | 4.0 | 2.3     | 18  |
| Aircraft/Aviation | Before 2010  | 4.1  | 7.5 | 9.6    | 10.4  | 13.3 | 21.2 | 5.7 | 3.8     | 33  |
|                   | Since 2010 * | 6.2  | 6.3 | 7.0 *  | 7.0 * | 7.5  | 7.7  | 1.2 | 0.5     | 4   |
| C4I               | Before 2010  | 1.1  | 3.3 | 7.8    | 7.0   | 10.1 | 12.5 | 6.8 | 3.6     | 14  |
|                   | Since 2010   |      |     | 5.6    | 5.6   |      |      |     | 0.0     | 1   |
| Electronic        | Before 2010  |      |     | 8.8    | 8.8   |      |      |     | 0.0     | 1   |
|                   | Since 2010   |      |     |        |       |      |      |     |         | 0   |
| Ground Systems    | Before 2010  | 3.0  | 4.0 | 6.8    | 6.3   | 8.2  | 8.7  | 4.3 | 2.1     | 4   |
|                   | Since 2010   | 5.0  | 5.0 | 7.7    | 7.0   | 8.4  | 8.4  | 3.4 | 1.5     | 3   |
| Missile/Munitions | Before 2010  | 2.5  | 5.0 | 7.0    | 7.3   | 9.3  | 14.6 | 4.4 | 3.0     | 24  |
|                   | Since 2010   | 2.8  | 3.2 | 3.7    | 4.9   | 7.2  | 9.5  | 4.0 | 2.4     | 5   |
| Mission Support   | Before 2010  |      |     | 7.5    | 7.5   |      |      |     | 0.0     | 1   |
|                   | Since 2010   |      |     |        |       |      |      |     |         | 0   |
| Ship/Maritime     | Before 2010  | 4.8  | 5.9 | 9.2    | 8.8   | 11.4 | 11.8 | 5.5 | 2.5     | 8   |
|                   | Since 2010 * | 0.8  | 2.0 | 3.9 *  | 4.3 * | 6.9  | 9.2  | 4.9 | 2.7     | 5   |
| Space             | Before 2010  | 5.2  | 6.6 | 9.7    | 9.8   | 13.0 | 13.8 | 6.4 | 3.0     | 8   |
|                   | Since 2010   |      |     |        |       |      |      |     |         | 0   |

\* Statistically significant difference.

NOTES: Includes only actuals cycle times in years from achieved MS B to achieved IOC (or equivalent). The cycle times Overall, for Aircraft/Aviation, and for Ship/Maritime capability areas were statistically lower since January 2010 than before (Mann-Whitney Test for unpaired data; p-values were 0.0008, 0.05, and 0.01, respectively). Differences for other capability areas were not statistically significant. The one outlier shown was in the Aircraft/Aviation area before 2010. If the outlier at 21.2 years was removed, the max. would be 16.3 years.


Figure 2. Time from MDAP Initiation at MS B to Achieved IOC (Before and Since January 2010)

## MDAPs Initiated at MS C

Very few MDAPs started at MS C—either before or since January 2010 (11 before and 6 since). As shown in Figure 3, there were no statistical differences between the two populations.



While the median and averages were similar, the variation was higher in the population before 2010 (again, with a very small population size). Thus, initiating MDAPs at MS C (i.e., with relatively mature technologies not requiring development) may be delivering at about the same speed, but caution is warranted because of the very small number of MDAPs initiating at MS C.




|             | Min. | Q1  | Median | Avg. | Q3  | Max. | IQR | St.Dev. | N = |
|-------------|------|-----|--------|------|-----|------|-----|---------|-----|
| Before 2010 | 0.7  | 0.8 | 2.7    | 3.8  | 7.2 | 8.7  | 6.4 | 2.9     | 11  |
| Since 2010  | 2.3  | 2.6 | 3.5    | 3.5  | 4.2 | 5.0  | 1.6 | 0.9     | 6   |
| Overall     | 0.7  | 1.5 | 3.2    | 3.5  | 5.3 | 8.7  | 3.8 | 2.4     | 17  |

**Figure 3. Time from MDAP Initiation at MS C to Achieved IOC (Before and Since January 2010)**

### MDAPs Initiated at Either MS B or C

Finally, when combining MDAPs initiated at either MS B or C, we find the cycle times slightly lower than viewing just the MS B MDAPs, probably because the MS C MDAPs tend to be faster to IOC as a population. **Figure 4** plots the results. The average cycle time dropped from 8.1 to 5.1 years overall. Aircraft/Aviation and Ships/Maritime MDAPs continued to be faster since 2010 (dropping from 9.9 to 5.6 years on average for the former and 7.3 to 4.3 years for the latter). Also, there were three Missile/Munition MDAPs that started with MS C, so that capability area was also statistically faster, dropping from 7.3 to 4.3 years on average. The remaining capability areas showed no statistically significant differences (again, with very small population sizes probably limiting the statistical tests).





\* Statistically significant difference.

NOTES: Includes only actuals cycle times in years from achieved MS B or C (if the program skipped MS B and initiated at MS C) to achieved IOC (or equivalent). The cycle times *Overall* and for *Aircraft/Aviation*, *Missiles/Munitions*, and *Ship/Maritime* capability areas were statistically lower since January 2010 than before (Mann-Whitney Test for unpaired data; p-values were 0.0002, 0.009, 0.01, and 0.01, respectively). Differences for other capability areas were not statistically significant. The one outlier shown was in the *Aircraft/Aviation* area before 2010. If the outlier at 21.2 years was removed, the max. would be 16.3 years.

**Figure 4. Time from MDAP Initiation at MS B or C to Achieved IOC (Before and Since January 2010)**


## Discussion

*Figure 5* provides a summary view of the three cases: MS B, MS C, and MS B/C MDAPs (independent of the capability portfolio details in earlier figures). Again, statistically significant



improvements in speed are seen in the first and third cases but not in the smaller but already faster MS C cases.

These results show that efforts to speed acquisition are having a significant, measurable effect—albeit with a relatively small population size and some caution about the types of MDAP systems before and after these time periods. For example, the complexity of the MDAPs initiated since 2010 may be a factor. Still, these results show that efforts to speed acquisition are having a significant effect even for the Department's largest and most complex weapon systems: MDAPs.



\* Statistically significant difference.

Figure 5. Time from MDAP Initiation to Achieved IOC (Before and Since January 2010)

It is also interesting that the central tendency of MDAPs initiated since 2010 achieve IOC in about the same time as Middle Tier program are supposed to deliver: 5 years (Under Secretary of Defense for Acquisition and Sustainment, 2020, pp. 4–5).

Software Acquisition Pathway seek delivering minimum viable products (or minimum viable capability releases) in 1 year or less (i.e., well under 5 years), but most of these are not at the size and complexity of an MDAP (Under Secretary of Defense for Acquisition and Sustainment, 2020, pp. 3, 16).



## Suggestions for Future Research

It will be useful to see if this trend in speed to IOC for MDAPs continues. Also, further research is warranted to compare the characteristics of the MDAPs before and after 2010 to seek any potential drivers; for example, are these size and complexity effects, or are MDAPs using best practices of utilizing mature technology or Agile development methods to speed time to fielding. In particular, understanding what may be traded for speed (e.g., system performance, resilience, or cybersecurity) would be useful to inform future decisions in how to more rapidly deliver capabilities to warfighters and other operators.

## Summary

Surprisingly (given general beliefs about the speed of defense acquisition), our analysis found that recent MDAPs reached IOC statistically faster than historical ones. The average and medians are now about 5 years instead of the prior 7–8 years. Also, MDAPs that leveraged mature technologies (a best practice for speed) tended to have reached IOC in about 3.5 years on average and thus could be a pattern for wider use.

While passing statistical tests, some caution is warranted because the population sizes are relatively small (especially for individual capability areas), and the mix in types and complexity of the MDAPs since 2010 may differ from the population of MDAPs starting before 2010. Nevertheless, these results show that efforts to speed acquisition are having a significant effect—even for the Department’s largest and most complex weapon systems—and speed can be increased (albeit through various practices and tradeoffs).

## Acronyms and Abbreviations

|         |                                        |
|---------|----------------------------------------|
| AIRC    | Acquisition Innovation Research Center |
| AT&L    | Acquisition, Technology, and Logistics |
| DoD     | Department of Defense                  |
| CY      | Calendar Year                          |
| IOC     | Initial Operating Capability           |
| IQR     | Interquartile Range                    |
| Max.    | Maximum                                |
| MDAP    | Major Defense Acquisition Program      |
| Min.    | Minimum                                |
| MS      | Milestone                              |
| MSAR    | Modernized Selected Acquisition Report |
| N       | Population Size                        |
| Q       | Quartile                               |
| Q1      | First Quartile                         |
| Q3      | Third Quartile                         |
| SAR     | Selected Acquisition Report            |
| SERC    | Systems Engineering Research Center    |
| St.Dev. | Standard Deviation                     |



|      |                                       |
|------|---------------------------------------|
| UARC | University-Affiliated Research Center |
| UCLA | University of California, Los Angeles |
| U.S. | United States                         |
| USD  | Under Secretary of Defense            |

## Acknowledgements

We would like to thank Laura Cooper for her sponsorship and support of our analytical effort and Pete Modigliani and Brian Joseph for their encouragement, support, and comments. Timothy Rodriguez provided data extracts from Advana and scripted initial calculations of time to IOC, which were cleaned by the author. Doug Buettner provided useful comments on early drafts of the analysis. Nickee Abbott, Daniel Hettema, Gwen Grewe, and Tara Kelly facilitated security reviews and public release of this report. Kara Pepe provided overall leadership of the research task award.

This material is based upon work supported, in whole or in part, by the U.S. Department of Defense through the Office of the Under Secretary of Defense for Acquisition and Sustainment (OUSD(A&S)) and the Office of the Under Secretary of Defense for Research and Engineering (OUSD(R&E)) under contract # HQ0034-19-D-0003, Task Order 0198, and contract # HQ0034-24-D-0023, Task Order 048. The views, findings, and conclusions in this document are solely those of the authors and do not necessarily reflect the views or positions of the U.S. Government (including the Department of Defense [DoD] or any other government personnel) or the Stevens Institute of Technology.

**Disclaimer:** Cleared for Public Release. Prepublication Copy. Anton, Philip S. (forthcoming). *23<sup>rd</sup> Annual Acquisition Research Symposium & Innovation Summit*, Naval Postgraduate School. Virtual. May 6–7, 2026. <https://nps.edu/web/acqnresearch/symposium>

## References

- Defense Acquisition University. (2026). *DAU glossary of defense acquisition acronyms and terms*. <https://www.dau.edu/glossary>
- Exec. Order No. 14265 (2025). <https://www.federalregister.gov/documents/2025/04/15/2025-06461/modernizing-defense-acquisitions-and-spurring-innovation-in-the-defense-industrial-base>
- Secretary of Defense. (2017, December 18). *2017 national defense strategy of the United States*. <https://trumpwhitehouse.archives.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905.pdf>
- Secretary of Defense. (2022, October 27). *2022 national defense strategy of the United States*. <https://media.defense.gov/2022/Oct/27/2003103845/-1/-1/1/2022-national-defense-strategy-npr-mdr.pdf>
- Secretary of War. (2025a, November 7). Transforming the Defense Acquisition System into the Warfighting Acquisition System to accelerate fielding of urgently needed capabilities to our warriors [Memorandum]. <https://media.defense.gov/2025/Nov/10/2003819439/-1/-1/1/transforming-the-defense-acquisition-system-into-the-warfighting-acquisition-system-to-accelerate-fielding-of-urgently-needed-capabilities-to-our-warriors.pdf>
- Secretary of War. (2025b, November 7). *Acquisition transformation strategy*. <https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/acquisition-transformation-strategy.pdf>



- Secretary of War. (2026, January 23). *2026 national defense strategy*.  
<https://media.defense.gov/2026/Jan/23/2003864773/-1/-1/0/2026-national-defense-strategy.pdf>
- Under Secretary of Defense for Acquisition and Sustainment. (2020). *Operation of the software acquisition pathway* (DoD Instruction 5000.87).  
<https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500087p.PDF>
- Under Secretary of Defense for Acquisition and Sustainment. (2024, November 25). *Operation of the middle tier of acquisition* (DoD Instruction 5000.80).  
<https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500080p.PDF?rnd=60820>
- Under Secretary of Defense for Acquisition, Technology, and Logistics. (2014, September 19). *Better buying power 3.0: Interim release*. [https://csis-website-prod.s3.amazonaws.com/s3fs-public/legacy\\_files/files/attachments/140919\\_BBP\\_3\\_0\\_InterimReleaseMaterials.pdf](https://csis-website-prod.s3.amazonaws.com/s3fs-public/legacy_files/files/attachments/140919_BBP_3_0_InterimReleaseMaterials.pdf)





# **Weapon Systems Annual Assessment: DoD Leaders Should Ensure That Newer Programs Are Structured for Speed and Innovation**

**Erin Carson**—is an Assistant Director with the U.S. Government Accountability Office, specializing in DoD weapon systems acquisition. She has worked in this area for nearly 15 years, reporting and making recommendations to improve DoD weapon systems acquisition outcomes. She is currently managing the GAO's annual weapon systems assessments, as well as reviews of the DoD's efforts to implement digital engineering. Erin holds a Master of Public Policy degree from the University of Denver's Josef Korbel School of International Studies, a Master of Arts degree from the University of Denver, and a Bachelor of Science degree from the University of Colorado. [carsone@gao.gov]

## **Abstract**

Despite recent reforms, the Department of Defense (DOD) remains plagued by escalating costs, prolonged development cycles, and structural inefficiencies that impede its ability to acquire and deploy innovative technologies with speed. GAO found that DOD plans to invest nearly \$2.4 trillion to develop and acquire its costliest weapon programs. Yet the average expected time for major defense acquisition programs to provide the warfighter with even an initial capability is almost 12 years from the program's start—a time frame incompatible with meeting emerging threats. DOD remains deeply entrenched in a traditional linear acquisition structure—characterized by rigid, sequential processes—that has proven inadequate in adapting to evolving threats and integrating emerging innovation. In a linear acquisition, the cost, schedule, and performance baselines are fixed early. Thus, programs develop weapon systems to meet fixed requirements that were set years in advance. This approach risks delivering a system—sometimes decades later—that is already obsolete. In contrast, leading companies use iterative cycles to design, validate, and deliver complex products with speed. Activities in these iterative cycles often overlap as the design undergoes continuous user engagement and testing, which allows the product to get to market quickly. This presentation discusses findings from GAO's 23rd annual weapon system assessment, including characteristics and performance of 106 of DOD's costliest weapon programs and findings from selected programs' implementation of leading practices for product development, as well as modern software and cybersecurity practices.

## **Background**

The DoD plans to invest nearly \$2.4 trillion to develop and acquire its costliest weapon programs, but it continues to struggle with delivering timely and effective solutions to the warfighter. Weapon systems are more complex and software-driven than ever before. The DoD implemented recent reforms intended to lead to faster results, but slow, linear development approaches persist.

This report, the GAO's 23rd annual assessment, responds to a provision Congress included in statute for the GAO to annually review selected DoD acquisition programs and efforts. It assesses the characteristics and performance of 106 of the DoD's costliest weapon programs. It further assesses selected programs for schedule performance, implementation of leading practices for product development, and approach to software development and cybersecurity practices.

## **DoD Plans to Invest Nearly \$2.4 trillion in Its Weapon Portfolio**

The weapon systems portfolio we assessed continues to grow both in cost and number of programs. It consists of 106 programs: 79 major defense acquisition programs (MDAPs), 20 middle-tier of acquisition (MTA) programs, and seven future major weapon acquisitions not



currently on an adaptive acquisition framework (AAF) pathway. The following sections provide an overview of the DoD's MDAP and MTA program portfolios.

## **DOD MDAP Portfolio**

**MDAP Cost.** The estimated cost change of programs that were also in the DoD's prior-year MDAP portfolio increased by \$119.9 billion. Among the 76 MDAPs for which cost data was available, half of the programs (38 of 76) reported cost increases totaling \$133.4 billion. The DoD plans to invest \$792.2 billion to develop and produce the subset of 36 MDAPs that, as of January 2025, were generally between the start of development and the early stages of production—over a third of the DoD's total estimated MDAP costs. Combined total estimates increased by \$49.3 billion for 30 MDAPs also included in the prior year's report. The Air Force's Sentinel missile program accounted for over \$36 billion (73%) of this increase.

**MDAP Schedule.** The expected time for MDAPs in the DoD's portfolio to provide an initial capability is almost 12 years from the program's start—an average that includes MDAPs that began as MTAs. Since our prior-year report, six programs in our review delayed their expected initial operational capability (IOC) by 12 months or more. These six programs have also reported delays previously. Among the 30 MDAPs overall that have yet to achieve IOC, 24 have delayed their expected time frames since first full estimate.

## **DOD MTA Portfolio**

**MTA Cost.** The DoD plans to invest at least \$44.5 billion across 20 of its largest MTA programs. Combined costs for 14 MTA programs included in both this report and the prior year's assessment trended up slightly by \$748 million (3%). MTA costs include ongoing efforts only and do not include any further investments the DoD may make to develop or acquire a capability after the current MTA effort concludes.

**MTA Technology Development.** MTA programs that initiated with immature critical technologies reported limited development progress. The 10 MTA programs with critical technologies that we reviewed in both this report and the prior year's assessments varied in their progress toward maturing them. Several programs that initiated with immature technologies still need significant work before reaching the end of their MTA effort and delivering initial capability.


**MTA Schedule.** MTA programs reported a wide range of time frames for delivering initial capabilities. We found that more programs are planning to transition to the MCA pathway at production start, operations and sustainment, or to an MTA rapid fielding effort than in prior years. However, of the 15 programs that reported plans to achieve an IOC, six programs still plan for further development for multiple years after ending the current MTA efforts before delivering capability.

## **Programs Are Missing Opportunities to Deliver Capability With Speed**

Opportunities exist for future major weapon acquisitions that have yet to start on an adaptive acquisition pathway to leverage leading practices during the earliest stages of the program—before they become locked into rigid requirements, budgets, and development approaches. These future programs reported that they intended to incorporate leading practices generally at levels at or below the levels reported by current MDAPs or MTAs. This is because decision-makers in the DoD and across the military services do not take steps to ensure that future programs include leading practices (discussed below). Incorporating leading practices prior to formally starting a new program can help programs take full advantage of the efficiencies they provide.



Most programs the GAO reviewed do not fully implement leading practices in concert to achieve efficiencies (Figure 1). For example, most programs reported using a modular open systems approach—generally required by statute—which allows them to easily add or replace weapon parts over time. Few, however, reported plans to establish a minimum viable product (an initial set of capabilities that can be iterated upon), use digital twinning (a virtual representation of a physical product), or use digital threads (real-time data to inform decision-making). These practices are most effective when they are used together as part of an iterative approach to product development.



Source: GAO analysis of programs' questionnaire responses. | GAO-25-107569

Most programs we reviewed do not fully implement leading practices, including future efforts that are newer and have opportunities to do so we are recommending that the DoD leadership take steps to ensure that future major weapon acquisition programs include leading practices for product development during the earliest stages of the programs. The DoD concurred with the recommendations.

*This is an excerpt from a full-length report. See GAO-25-107569 for additional details, including additional report contributors: [gao.gov/products/gao-25-107569](https://www.gao.gov/products/gao-25-107569) (report landing page); [gao.gov/assets/gao-25-107569.pdf](https://www.gao.gov/assets/gao-25-107569.pdf) (PDF of report).*



## PANEL 13. FROM DEMO TO DETERRENCE: ACCELERATING SURGE CAPACITY AND FLEET INTEGRATION

Wednesday, May 6, 2026, 1615 – 1730 ET (1315 - 1430 PT)

### Panel Summary:

Success in modern strategic competition is determined by the ability to shape the environment before conflict and adapt at the speed of battle. This panel explores breakthrough frameworks for enhancing warfighting readiness through adaptive contracting, advanced manufacturing for legacy munitions, and the rapid 30-day fielding of critical maintenance technologies. By integrating AI-enabled infrastructure investment tools with high-level advocacy and creative transition authorities, these research initiatives provide a blueprint for overcoming programmatic barriers and delivering resilient, integrated capabilities to the fleet.

**Chair: Dr. Dave Lewis, VADM USN (ret.),** Professor of the Practice, Naval Postgraduate School, former Director of the Defense Contract Management Agency (DCMA)

### Presenters:

**Pathfinding Innovation for Legacy Munitions** – *Gregory Sanders, Joint Production Accelerator Cell (JPAC)*

**Generating Adaptive Capability from Constraints: Aligning Operational Dependency, Requirements, and Competitive Procurement to Sustain Advantage in Maritime Expeditionary Operations** – *Lieutenant Commander Adam Pierce, Navy Expeditionary Warfighting Development Center*

**Generating Adaptive Capability from Constraints: Aligning Operational Dependency, Requirements, and Competitive Procurement to Sustain Advantage in Maritime Expeditionary Operations** – *Matthew Cole, Naval Surface Warfare Center, Port Hueneme Division*



**Dr. Dave Lewis, VADM USN (ret.)**—is a retired United States Navy Surface Warfare Officer and Engineering Duty Officer. He currently serves as a Professor of the Practice at the Naval Postgraduate School, Monterey, CA as an intermittent employee. He works full time in the Defense Industry.

His leadership assignments included duty as the Aegis Shipbuilding Program Manager, Program Executive Officer, Ships, Commander, Space and Naval Warfare Systems Command, and prior to his retirement, as Director of the Defense Contract Management Agency.

His afloat tours were in USS Spruance (DD-963), USS Biddle (CG-34) and USS Ticonderoga (CG-47), all in combat systems and communications assignments.

David Lewis earned bachelor's and master's degrees in computer science and a certificate of completion in Weapons Engineering. He recently completed his Ed.D in Leadership and Learning in Organizations. He is a graduate of the Naval War College, Vanderbilt Peabody College, the Naval Postgraduate School and the University of Nebraska.

He is a lifetime Member of the US Naval Institute, the American Society of Naval Engineers and the Surface Navy Association and has published several articles on shipbuilding and warship engineering topics.



# Pathfinding Innovation for Legacy Munitions

**Laura Leichus**—is a Strategy & Innovation Analyst at the Office of the Undersecretary of War for Acquisition and Sustainment and a Contractor with Booz Allen Hamilton. In her current role, she works to help facilitate surging munitions production through coproduction agreements and advanced manufacturing approaches.

**Greg Sanders** —is a Senior Policy Analyst at the Office of the Undersecretary of War for Acquisition and Sustainment and the Center for Strategic and International Studies. He also works to surge munitions production through coproduction and other means, with a particular interest in Modular Open System Approaches.

## Abstract

Advanced manufacturing encompasses the innovation of improved methods for manufacturing existing products, such as components and parts, as well as the production of new products enabled by advanced technologies. Advanced manufacturing approaches, include robotic technologies, digital engineering, sensing, advanced materials, and 3D printing are reimagining manufacturing. The Pathfinding Innovation for Legacy Munitions Project seeks to understand the landscape of advanced manufacturing approaches and to identify processes that could be applied to legacy munitions to enable surging production of major weapons systems. In this paper, legacy munitions refer to munitions that are currently in production and have been deployed. Currently, many legacy munitions are manufactured using production processes that predate proven commercial approaches. Newer commercial approaches have the potential to improve the cost-effectiveness and efficiency of munitions production.

## Executive Summary

A range of challenges impede implementing advanced manufacturing in legacy munitions but present efforts to ramp production mean that government and industry have reason to consider new solutions to pressing problems. Identifying which of these approaches can be economically applied to legacy munitions enables proven capabilities to reach the warfighter faster.

The projects findings are based on 18 interviews with acquisition and sustainment practitioners and subject matter experts (SME), including the Manufacturing Innovation Institutes (MIIs) Manufacturing x Digital (MxD) and Leading Innovations for Tomorrow (LIFT), and the Advanced Robotics for Manufacturing Institute (ARM). These interviews were used to determine where advanced manufacturing is being applied or is mature enough to apply to produce legacy munitions.

While examples of these technologies successfully applied to legacy munitions exist, many obstacles remain to their incorporation. These obstacles fall into the following categories:





**Table 7. Potential Reforms and Best Practices for Implementing Advanced Manufacturing for Legacy Munitions**

| <u>Producer Financial Incentives</u>                                                                                                                                                                                                                                         | <u>Departmental Process and Regulations</u>                                                                                                                                                                                                                                                       | <u>Culture and Workforce Constraints</u>                                                                                                                                     |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>• Demand Signal and Boom and Bust</li> <li>• Limits of Incremental Adoption</li> <li>• Narrow Time Windows for Economical Adoption</li> <li>• Barriers to Competition</li> <li>• Intellectual Property and Business Models</li> </ul> | <ul style="list-style-type: none"> <li>• Qualification and Testing Challenges</li> <li>• Safety Concerns</li> <li>• Contract, Requirement, and Design Inertia</li> <li>• Appropriations Process</li> <li>• Contract, Requirement, and Design Inertia</li> <li>• Appropriations Process</li> </ul> | <ul style="list-style-type: none"> <li>• Resistance to Change</li> <li>• The First Mover Disadvantage for Programs</li> <li>• Workforce Requirements and Concerns</li> </ul> |

The project found that advanced manufacturing has mature approaches in the spaces of automation, digital engineering/sensing, and advanced materials. Robotics reduced the need for humans to handle repetitive or dangerous tasks. Digital twins and sensing technologies have reduced reliance on operators to determine the readiness of mixtures, as demonstrated at U.S. Army Combat Capabilities Development Command (DEVCOM) Armaments Center (AC). Lastly, the Advanced Metallic Production and Processing (AMPP) Center produced new metallics and alloys can performance enhancements. In addition, additive manufacturing and flexible manufacturing are promising future approaches.

**Table 8. Mature Advanced Manufacturing Processes Showing Promise for Legacy Munitions**

| <b>Type of Advanced Manufacturing</b> | <b>Facility / Location</b>                 | <b>Government Organizers</b>          | <b>Details</b>                                                                     |
|---------------------------------------|--------------------------------------------|---------------------------------------|------------------------------------------------------------------------------------|
| Automation                            | Troy, AL                                   | ARM Institute                         | Robotic sanding and finishing solution                                             |
|                                       | Wharton, NJ                                | U.S. Army DEVCOM AC                   | Container for handling energetics, robot code, and a supervisory system controller |
|                                       | Wharton, NJ                                | U.S. Army DEVCOM AC and ARM Institute | Robot scoops and weighs powder                                                     |
|                                       | Marion, IL, Hampton, AR, and Coachella, CA | DPAP                                  | Robotic stations for production line                                               |
|                                       | Lincoln, NE                                | DPAP                                  | Nondestructive inspection of SRM nozzles.                                          |
|                                       | Wharton, NJ                                | Army ManTech                          | Slurry loading                                                                     |
|                                       | Camden, AR, Orange, VA, and Huntsville, AL | DPAP                                  | Robots move SRM cases and components                                               |
|                                       | Xenia, OH                                  | DPAP                                  | Producing SRM cases using additive manufacturing                                   |
| Digital Engineering and Sensing       | Wharton, NJ                                | MxD and U.S. Army DEVCOM AC           | Machine-agnostic sensor kit that gathers data for analysis                         |
|                                       | Wharton, NJ                                | U.S. Army DEVCOM AC and MxD           | Digital twin for melt-pour process                                                 |
| Advanced Materials                    | Detroit, MI                                | LIFT AMPP                             | Creating custom alloys and feedstocks for metals                                   |



These examples show that advanced manufacturing has proven benefits that could transform parts of legacy munitions production. The paper identifies reforms and best practices to ease introduction.

**Table 9 Potential Reforms and Best practices for Implementing Advanced Manufacturing for Legacy Munitions**

| <u>Producer Financial Incentives</u>                                                                                                                                                                                                                                                                                                   | <u>Departmental Process and Regulation</u>                                                                                                                                                                                                                                                                | <u>Culture and Workforce Constraints</u>                                                                                                                     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>• Create a steady demand signal</li> <li>• Invest in research lab efforts</li> <li>• Derisk and pilot applications</li> <li>• Encourage MII, SRL, and consortium collaboration</li> <li>• Use OTAs, fixed price contracting, and multi-year contracting</li> <li>• Distribute orders</li> </ul> | <ul style="list-style-type: none"> <li>• Share testing data</li> <li>• Change the qualification process</li> <li>• Accept manufacturer and supplier data</li> <li>• Accept sensor data and digital twins</li> <li>• Use OTAs to support prototyping, partnerships, and distributed contracting</li> </ul> | <ul style="list-style-type: none"> <li>• Educate Communities</li> <li>• Data Driven Decision-making</li> <li>• Training and Workforce Development</li> </ul> |

JPAC is engaging partners on promising approaches and launching a pilot to surge production and expand use of advanced manufacturing approaches by the defense industrial base. This project will focus on fostering proactive failure analysis detection and mitigation.

## The Introduction

The United States Defense Industrial Base (DIB) is unable to surge munitions production to build weapon stockpiles within an acceptable timeframe. Business-as-usual munition production rates pose strategic and operational risks. The study supports the Joint Production Accelerator Cell's (JPAC) mission to determine how the United States can surge the production of critical munitions and ensure that the United States is prepared to counter the pacing challenge and evolving threats within an acceptable timeframe. Business as usual munition production rates lead to strategic and operational risks.

One way that the Department of Defense (DoD) could accelerate production is by applying advanced manufacturing approaches to the production of already established weapons systems, known as legacy systems or legacy munitions. Examples of this include using automation in place of manual tasks, using sensing data to collect and organize data to optimize manufacturing in real-time, and using advanced materials to support parts production and address supply chain vulnerability. These legacy munitions are already in use and have well-established production lines and supply chains, reliant on older technologies and processes. Changing manufacturing approaches for legacy munitions is expensive and time-consuming. However, applying innovative technologies, such as advanced manufacturing approaches, could make legacy munitions manufacturing cheaper, faster, and more agile, unlocking the potential to surge production capability. In the case of energetics, advanced manufacturing techniques can enhance worker safety by removing as much human exposure to energetics materials as possible.

## Study Approach

The Pathfinding Innovation for Legacy Munitions project is studying how and when advanced manufacturing approaches can be efficiently applied to improve production capacity for the legacy munitions enterprise. The project employed five overarching research questions to conduct 18 interviews with Manufacturing Innovation Institutes (MIIs), outside researchers, and experts within the Department to identify the challenges the DIB faces in implementing



advanced manufacturing. Interviews were with Program Managers, Engineers, Technical Advisors, Special Assistants, Directors, Chief Executive Officers, and Chief Technology Officers, amongst other experts. These interviews identified challenges with producer financial incentives, departmental processes and regulations, and culture and workforce constraints.

These research questions are as follows:

- What advanced manufacturing approaches or commercial technologies are currently being applied to enhance the producibility of legacy munitions?
- What advanced manufacturing approaches or commercial technologies are relevant and mature, but not yet applied to legacy munitions?
- Why are advanced manufacturing approaches not more widely used in the DoD? What could be done to change that?
- What challenges are industrial partners that employ advanced manufacturing approaches facing when working with the DoD or legacy munitions?
- What are the limitations for advanced manufacturing approaches in terms of producibility?
  - Is there a need for more testing to mature the technologies?
  - Are the advanced manufacturing approaches difficult to qualify?
  - Is cost a limiting factor?
  - Is there a need to incorporate advanced manufacturing approaches early to avoid the need to modify processes or designs? Is this causing problems with producibility or manufacturing?

## **The Obstacles**

There are many reasons that advanced manufacturing is not more widely applied and integrated throughout the Department. These issues, summarized in Table 7, include a lack of incentives to adopt new methods, inertia in current processes, and cultural resistance from some parts of the munition enterprise and IB. Incorporating advanced manufacturing can introduce risks to balance-sheets production lines, and performance and safety requirements and may encounter cultural resistance when being introduced into the manufacturing process for legacy systems. Resources and effort more easily flow to new systems that can incorporate these approaches at the start. However, as subsequent sections will discuss, JPAC has found that the DoD misses essential opportunities due to a lack of integration of promising emerging technologies.

### **Producer Financial Incentives**

#### ***Demand Signal and Boom and Bust***

Legacy munitions, by definition, have existing production lines, and those production lines were scoped based on industry and program expectations of current and projected demand from the United States and licensed exports. Unless there are deactivated and preserved tools, lines, or facilities, changing or expanding scope using traditional or advanced manufacturing methods requires major capital investments and does not address potential workforce constraints. Meanwhile, for a sole source provider, demand backlogs often represent revenue pushed into the future rather than foregone revenue. As a result, it is not economical to expand production facilities to meet a temporary spike in demand that will vanish after stockpiles have been refilled.



Without robust forecast or assurances that there will be demand for the legacy munition, industry cannot be confident that investments in advanced manufacturing that are specific to a certain munition will be profitable. Slack capacity is sometimes intentionally maintained in government-owned facilities, but this is a choice that must compete with other priorities, such as facilities that could better support higher performance or affordable mass munitions now under development. The problem for both the private and publicly owned facilities is that munitions manufacturing is prone to boom-and-bust cycles. Guaranteed government interest through multi-year contracting or minimal purchase guarantees are two statutory tools to manufacturers and suppliers that there will be a sustained need for the approach, munition, or component.

### **Limits of Incremental Adoption**

Achieving the full benefits of advanced manufacturing often requires investments in infrastructure, new supply chains, professional staff time, and data. Cultural change also undercuts piecemeal adoption of advanced manufacturing. Similarly, established public and private sector suppliers also often lack the know-how to modernize their outdated infrastructure. The culture and expertise challenges are discussed later in this paper under *Culture and Workforce Constraints*. Ideally, advanced manufacturing approaches involve not just capital expenditure but also infrastructure that supports the process. Physical infrastructure often needs to be transformed to support advanced manufacturing techniques.

For example, robots may be able to handle materials or work enable agile automation but rely on sufficient space or leveraging digital engineering. Digital engineering is an integrated approach that uses digital models, data, and advanced computing technologies as a continuum across disciplines to support the entire life cycle of a system, from concept through disposal. For digital engineering to be employed, there needs to be a digital foundation and digital framework to provide factories with integrated technological infrastructure, systems, and data management frameworks that enable digital manufacturing environments and data-sharing. A key challenge for legacy munitions is that if engineering data was not created in a digital format during initial design or subsequent modifications, the extent of potential digital engineering is limited. Even if the data is available, the difficulties associated with integrating a digital foundation mean that many discussions about employing advanced methods are only had during the construction of a facility or ordnance plant. These locations are reluctant to make costly updates to their outdated infrastructure, which makes it challenging to integrate new approaches.

A key early step to a digital foundation is networking. A networked advanced manufacturing approach significantly increases speed and reduces cost through enhanced efficiency, real-time data flow, automation, and improved decision-making across the entire product life cycle. Factories need execution software and network communications to ensure data can be collected in real time and to provide insights that enable quick adjustments to avoid latency, lags, and machinery downtime. Many current production sites have machines or computers that have no network interfaces or are otherwise not connected to wired or even wireless networks, referred to as “air gapping.” Air gapping also makes it difficult to apply security patches to software and make technological changes. Similarly, air gapping prevents machinery from recording and reporting data in real-time, and can lead to data being recorded on paper, which hampers analysis and process control. Manual data collection is often cumbersome, which prevents it from occurring or being employed. This isolation of machines or computers may be the result of industrial design approaches that pre-dated digital engineering or specific cybersecurity or safety concerns. Production lines with isolated machines face significant friction when employing digital engineering. One interviewee cited a successful demonstration at Rock Island of 3D printing an antiquated part. Commercial experts in the audience noted that the machine was not connected to anything, which would make it difficult to further iterate on the success.



Deviating from established manufacturing plans has many cost implications beyond infrastructure upgrades. Manufacturers face costs in both time and money for establishing supply chains with suppliers that can support new technologies. New technologies may depend on materials and alloys that cannot be domestically sourced, are in high demand, or have long lead times, all of which hinder production. Often, incumbent vendors will purchase only from a limited, highly qualified supplier base.

These interconnected modernization requirements can present an additional challenge to government-owned facilities, referred to as the organic industrial base (OIB) and private sector facilities. Munition plants in the OIB typically lack the funds to modernize. For OIB facilities to support advanced manufacturing, they need updated infrastructure, hardware, and software. When the OIB facility is not the only potential source the DoD often prefers to rely on private industry rather than upgrading and investing in OIB facilities. Often the DoD is working with other partners, and those partners can make needed upgrades to their facilities. Likewise, the DoD may take approaches that spread funding across many facilities rather than making a difficult choice of which one or two to upgrade. This allows for incremental adoption but makes it difficult to make wholesale changes.

### **Narrow Time Windows for Economical Adoption**

For government or industry to invest in a new advanced manufacturing approach, it needs to make sense financially. Legacy munitions rely on existing facilities, production lines, and weapon system designs. The return on investment is most favorable if advanced manufacturing adoption aligns with upgrade opportunities in the life cycle of a facility, production line, or weapon system. Even for commercial manufacturers already using advanced manufacturing, manufacturing for the DoD accounts for only a small share of the business's revenue and therefore investing in re-engineering or re-tooling existing military-specific production lines may not yield a return on investment. Changing a manufacturing approach can also lead to tooling costs, which often requires a high initial investment. Tooling is often amortized over 10 to 15 years, and tools may take that long to be paid down. For both supply chains and capital equipment, manufacturers do not want to incur costs they did not budget for and that will not be reimbursed for. Specifically, the return on investment cost of adopting advanced manufacturing will often vary based on whether investments in current approaches or equipment are still depreciating.

Advanced manufacturing approaches do produce a range of benefits that should contribute to corporate revenue and help the OIB better perform its mission and meet contractual requirements. However, this business case is complicated for legacy munitions because to realize the full benefits of advanced manufacturing, manufacturers would need to commit to it early in the process, possibly before standing up facilities. Manufacturers may have passed the risky decision to adopt an approach that seemed immature when they first stood up a production line. Yet, if the approach is not adopted early, it can be hard to implement changes without disrupting production. Deciding to adopt an approach earlier is an easier decision if the approach has more advanced technological or manufacturing readiness levels. However, for advanced manufacturing approaches not yet applied in a military or commercial context, even mature advanced manufacturing approaches may require Research, Development, Testing, and Evaluation (RDT&E) spending or private internal R&D investments.

### **Barriers to Competition**

The commercial industry uses a range of advanced manufacturing approaches often because it provides an edge over competitors. Commercial technologies can be brought to market much faster and therefore speed of innovation is rewarded. The DoD frequently relies on sole-source production contracts, and those contractors are only incentivized to meet the





requirements. Even if a contractor captures additional revenue by expanding production, the necessary investments may lead to lower profit margins in a sole source environment. The difficulty in finding the correct scale to bundle contract orders is another problem. Large commercial companies are often uninterested in making the transition to becoming government contractors to win the quantities the government solicits because they are not large enough to yield a desirable return on investment. Small vendors are often unable to produce the required quantities and therefore do not bid.

### **Intellectual Property and Business Models**

The typical business model for DIB vendors rewards closely holding intellectual property (IP) and data rights rather than licensing production to other manufacturers, even if that other vendor may be capable of employing an advanced manufacturing approach. Incumbent vendors often hold proprietary designs and institutional knowledge that new entrants lack. New entrants seeking to become second source suppliers already face high upfront costs and regulatory requirements. If IP licensing is expensive or unavailable it may be impractical to become a second source for another company's product. Additionally, Original Equipment Manufacturers (OEMs) often have a vested interest in holding close information on the integration and final assembly of a product. This applies when the OEMs are subcontracting and is magnified when they manufacture the component themselves.

Unless the acquisition official in charge anticipated potential advanced manufacturing applications for a given legacy program, the DoD will typically not have clear rights to the Detailed Manufacturing or Process Data (DMPD) of the products it buys, which creates a dependence on the company that supplies the technology is purchased from. Even in those cases where the government has rights to DMPD, critical technical data might be lost or have contested usage rights. In these circumstances, government investments in process improvements may primarily benefit the recipient vendor's bottom line. If a vendor or the government seeks distributed production, concerns related to intellectual property and data rights would have to be worked out prior to working together. This may be costly for legacy munitions, as the IP Guidebook for DoD Acquisition warns, "buying the data, outside of competition and well into program execution, could be extremely expensive." A range of flexible IP arrangements is possible and discussed within the IP Guidebook, but this is another domain where applying advanced manufacturing is far easier for programs early in their life cycle or at very least undertaking incremental upgrades.

### **Departmental Process and Regulations**

#### **Qualification and Testing Challenges**

Vendors that want to certify advanced manufacturing processes to meet DoD standards must go through the qualification process.<sup>1</sup> Vendors wishing to adopt advanced manufacturing approaches must ensure the methods they intend to use are appropriately vetted, mature, and ready for use. New materials need to be verified and validated to ensure they are fit for the intended purpose and ready to be scaled up and can be used for broader industrial applications. These changes need to consider qualification and data acceptance standards. When starting up a production line using traditional or advanced approaches, the entire line, including both the process and the outcomes, must be certified. The process of qualifying limited rate and then full

---

<sup>1</sup> Commercial and defense standards include model based system engineering (MBSE), digital twins specifications, technological readiness levels (TRL), and manufacturing readiness levels (MRL), prototype development, and specific standards such as ISO/ASTM 52900, ASTM F2792. See the subsequent section on Qualification and Testing Challenges. For a discussion of one particular additive manufacturing approach, see the [DoD additive manufacturing strategy](#) for additional details.



rate production can be even harder for advanced manufacturing because there is less historical data to rely on, as discussed in the next sub-section. Moreover, legacy munitions already have production lines and processes that have undergone qualification. Industry finds the qualification process daunting because the process takes time and is expensive. OEMs also contribute to variation between programs due to their differing approaches. Each munition often requires an individual qualification, approval, and sign-off from a chief engineer.

Remaking legacy parts and components often will require requalification, which can be prohibitively expensive. Some components, such as energetics and fuzes, face particularly high burdens because new processes can introduce energy or change the chemical properties of the result. An end-product that matches one produced by a traditional method may not be accepted because there could be potential impurities or chemical changes not found in traditional methods. On the other hand, delaying qualification of a product may mean that a new entrant's investments in design, testing, and iterating to enhance producibility goes to waste because the item or subcomponent lacks key attributes to allow it to meet requirements. These burdens are greatest the first time an approach goes through development. Piloting for future adoption and cost reduction initiatives can ease adaptation in future programs, though adoption by legacy programs will still face higher burdens.

This process is lengthy due to extensive safety, performance, and reliability testing prior to delivering platforms to the public or military personnel. Because the qualification process is lengthy, manufacturers need to be prepared to outlay substantial funds for a potential benefit that will not be realized in the near term. Often, the expense does not yield a sufficient return on investment. Qualifiers may reject a technology, which makes it risky to invest in. Program offices for sustainment systems may pay to support company qualification and certification efforts, but only if they have sufficient O&M dollars. Proposals for qualification process reforms have been raised, which focus on allowing the use of prototyping, digital threads, artificial intelligence, or incremental assurance of capabilities.

### **Data Challenges**

Qualification requires supporting data. Developers of new technology often lack experience and data supporting the viability of their newly developed technology. Commercially proven technologies already have some data on hand, but that data may not address military specific concerns or the specific demands of government processes. As a result, vendors need to produce a statistically significant number of pieces or parts that can be tested to prove the validity of any process. This delays vendors' opportunity to earn a return on investment because the DoD requires this data when deciding whether to integrate or scale production of new technology, which amplifies these challenges. As part of the qualification process, validation often requires access to a controlled and intrinsically safe environments. Many existing facilities do not allow on-site testing and prototyping on their production lines. Thus, many vendors need to pay an upfront cost to contract with a government organization to make use of a testing range so their technology can undergo safety inspections. Small vendors and new entrants often do not have the funds to build a testing range, or it would not make sense for them to build one without confidence in a good return on investment or guaranteed purchase orders.

Vendors that manage to test their new process or end-item internally or generate testing data for the certification process, are often only rewarded with payment for all up rounds (AUR). If process controls are not in place, the DoD is unlikely to accept their data as allowable or to provide certification or qualification based on industry-provided data. This lack of acceptance of data disincentivizes vendors from investing in creating data or sharing data with the government.



## **Safety Concerns**

Vendors are also reluctant to move away from traditional manufacturing processes toward advanced manufacturing techniques due to safety concerns, particularly in energetics. Many munitions manufacturing steps are carefully controlled because they involve manual, labor-intensive, and dangerous work that can lead to injuries. Working with energetic materials is inherently risky, leading to stringent safety requirements for explosives, propellants, and related components such as motor cases. Vendors may be reluctant to change processes out of fear of undermining current safety measures.

Changes to processes implemented through automation are often met with resistance because they can create electromagnetic fields (EMF) and explosive dust, which contributes to the lengthy qualification process for new technologies involving energetics, as discussed above. Progress requires more knowledge about testing new technologies for use with energetics and using manufacturing equipment, convincing the engineering community to accept new and changing requirements, and then changing those requirements or processes. Quality control for a new process must validate the advanced manufacturing processes in a way that addresses safety concerns.

## **Contract, Requirement, and Design Inertia**

The Department primarily structures contracts to ensure stringent adherence to national security standards. The DoD offers some flexibility in contract requirements when warranted by offering waiver requests when suppliers do not meet certain contractual requirements.<sup>2</sup> The DoD is undertaking a modernization and commercial innovation adoption push and has contract approaches and acquisition policies to encourage new processes. However, this still means that the government and the IB often must do considerable work together to update contracts or requirements to allow for new processes, techniques, or surveillance approaches. For example, contracts often require an Authorization to Operate (ATO) before software can be deployed. This process can take months to years, depending on what stage of development the software is in, which may result in the software becoming obsolete by the time it is installed.

Engineering change proposals are a key mechanism to allow legacy munitions to implement advanced manufacturing. Engineering change requests will require time to get approved and may be rejected if funding is not available or if other changes are necessary, such as a need for new sources for materials. Avoiding delays to production or future contracts requires information to flow quickly and efficiently. Currently, the engineering change request approval process takes more time than implementing the change. Experts provided competing assessments as to whether it would be desirable to make this system more flexible, as it is built on years of learning from mistakes.

Legacy munitions have specific design constraints and performance requirements that are inconsistent with product changes. Changing design or individual performance requirements will have downstream implications, which will require updating guidance software, interface control documents, and performance requirements. Product designs often need to be updated to fully benefit from new manufacturing methods. Returning to the Rock Island 3D printing example raised under infrastructure requirements, additive manufacturing works best when parts are not simply manufactured using a new method but are redesigned to take advantage of 3D printing's versatility.

---

<sup>2</sup> Examples of flexibility under traditional contracting include adjusting payment terms, offering small business set-asides, and a willingness to offer economic pricing adjustments and requirements changes when appropriate.



## **Appropriations Process**

Implementing advanced manufacturing in legacy munitions can be an awkward fit for the defense appropriation process. Maturing new approaches is a natural fit for RDT&E, manufacturing of legacy munitions is the domain of procurement funding, upgrading the organic industrial base often involves military construction (MILCON) funding, addressing obsolescence issues is the domain of operations and maintenance (O&M). Funding account regulations, referred to as colors of money, mean that RDT&E, procurement, MILCON, and O&M funding have different rules and are segregated from one another. This can mean that any given funding line is often only relevant to part of the Producer Financial Incentives challenges discussed above. While this project is primarily concerned with the production of munitions, advanced manufacturing is also relevant to the sustaining of present and future stockpiles. Adopting advanced manufacturing can be a collective action problem since there are high upfront costs and the benefits are distributed across multiple communities within the government.

Fortunately, this problem is mitigated by the availability of Defense Production Act Title III and Industrial Base Analysis and Sustainment program funding. However, these mechanisms face multiple competing challenges and for the reasons discussed above and below, implementing advanced manufacturing for legacy munitions must compete with other urgent DIB resilience priorities.

## **Culture and Workforce Constraints**

### **Resistance to Change**

Often, both the government and industry are set in their ways and operating contractors for government owned contractor operator facilities are often negotiated in periods of ten or more years, with mandated requirements. Most of the time, the Department has an objective and vision for a product and contracts with an industrial partner that has a solution aligned with that vision. The Department chooses to contract with this partner based on the company's ability to provide a known solution to a problem. Because the Department has already bought into this solution, purchasers are reluctant to encourage change or tinkering with new ideas unless specific Contract Line-Item Numbers are funding the innovation. Federal Acquisition Regulations and the Defense Federal Acquisition Regulations allow for flexibility but slow the pace of cultural change. Another challenge is that the energetics community is risk-averse and has a low tolerance for newness due to concerns about safety, which are mentioned in the Safety Concerns section. Implementing something new often requires convincing stakeholders to be open to change.

### **The First Mover Disadvantage for Programs**

The first legacy program to adopt a given advanced manufacturing approach faces considerable challenges as discussed above but also creates spill-over benefits for other programs. However, a Program Manager's primary duty is to their own Program. Program Managers are responsible for the execution of approved production budgets within established guidelines. Their mission is to prioritize stability and predictability in established production lines to meet current procurement levels, control costs, maintain schedules, and deliver required performance. With lines often operating at high or near maximum utilization, disrupting these processes to experiment with new technologies introduces significant risks and potential liability from cost overruns, schedule delays, late deliveries, and potential disruptions to existing performance metrics. Consequently, Program Managers often prioritize proven processes and risk mitigation over potentially beneficial, but uncertain, innovations. Program Managers will not make changes if the current process is satisfying production rates. However, if there is a surge or an updated requirement that causes an increase to the required production rate, a Program Manager may be more interested in introducing new production techniques.



Unless the Prime is meeting its obligations, it will be challenging to get support for integrating advanced manufacturing from Program Offices, Program Executive Offices (PEOs), or individual Program Managers. Buy-in from Program Managers involved in production and quality control is necessary to begin implementing a new process. New ideas are most often met with suspicion, especially if they are not fully understood. Education is key, and there also needs to be trust in the latest technologies for innovation to be considered and approved. Conversely, when one legacy program adopts a novel advanced manufacturing approach, they can clear a path through the challenges discussed in the Departmental Process and Regulations section. Cross-program adoption is still costly and time consuming but Program Managers facing critical munition shortfalls today may prove to be the heralds of broader adoption.

### **Workforce Requirements and Concerns**

Implementing new manufacturing approaches can require finding new workers or retraining existing staff. Munitions manufacturing often involves low pay and poor or dangerous working conditions. The OIB and private sector already have difficulty finding workers willing to work on the factory floor. People with a 2-year college degree, a trade school degree, and/or specialized manufacturing knowledge are increasingly hard to find in today's world. The workforce may also need clearance to participate in classified manufacturing.

Advanced manufacturing raises productivity, but in the near term, it can compound labor shortages by requiring manufacturers to look for workers with specialized knowledge. Personnel operating these technologies need to be trained on new manufacturing systems before adoption. This training can include teaching workers how to program robots and operate machines employing advanced manufacturing, which can be a time-consuming and complex task. Shop floor personnel be prepared and empowered to make purchase decisions. Finally, post-processing work to iterate manufacturing approaches is an essential skill set for the manufacturing workforce.

Advanced manufacturing's specialized workforce requirements make it less appealing for established manufacturers to invest in technology. Also, it is difficult to spread technologies when specialized knowledge is limited to certain people in some geographic regions, which may not align with the location of existing facilities. It takes time to train more individuals to have the

### **Manufacturing Innovation Institutes**

To help mature manufacturing process that show promise, Office of the Secretary of Defense Manufacturing Technology Office (OSD ManTech) sponsors Manufacturing Innovation Institutes (MIIs), which are public-private partnerships that help to mature manufacturing processes and support required infrastructure. The MIIs connect technology, industry, and market sectors to enable the rapid transition of new defense technologies to the warfighter. More specifically, they advance research and technology through partnering with industry, establishing manufacturing ecosystems and regional hubs, and providing education and workforce development to the workforce.

There are nine MIIs in total including MxD, LIFT, and ARM Institute, which were the most relevant to this paper. These three MIIs aided in identifying examples of advanced manufacturing techniques being utilized to produce legacy munitions. In addition, to identifying examples, the MIIs were often involved in implementing many of the examples discussed below. MxD is a digital manufacturing and cybersecurity institute that works to accelerate the adoption of digital solutions and secure supply chains. LIFT is focused on supporting advanced manufacturing through technology, advanced materials, talent development, manufacturing processes, and systems engineering. ARM is focused on deploying robots and artificial intelligence to support implementing advanced manufacturing and developing the next generation workforce.





required skill sets if new workers not from the local area are needed.

For successful adoption of advanced manufacturing, a shared investment in workforce development is essential. Manufacturers need assurance that their workforce possesses the skills and knowledge to effectively support advanced processes, a significant challenge given the current skills gap. Reskilling initiatives are crucial but require time, resources, and a proactive approach to address potential worker concerns about job security and evolving roles. For example, those individuals who take on dangerous tasks are well compensated compared to laborers who are insulated from hazards by new production approaches. Emphasizing the opportunities for career advancement and skills enhancement can foster a more receptive environment for change.

## **Mature Advanced Manufacturing Processes Showing Promise**

Advanced manufacturing is already being applied to legacy munitions in a cases discussed below and summarized in Table 8. These existing examples are followed by a discussion of other promising manufacturing approaches.

### **Automation Technologies**

Automation is the use of technology in manufacturing processes to reduce the need for human intervention to perform in repetitive, complex, or dangerous tasks. Many of these tasks pose safety risks to the humans conducting them or can be completed by machines with greater consistency. Automation uses robotics, information technology, and computer-controlled machinery. By applying automation to existing processes, organizations can achieve enhanced efficiency, productivity, and flexibility. Automation achieves these benefits through mechanisms such as reducing waste through greater reliability, reducing touch labor, and leading to rapid reconfiguration of production lines. These mechanisms are particularly valuable for modernizing legacy systems, especially where precision and safety are paramount. Mature automation technologies show promise for enhancing the producibility of legacy munitions.

Implementing robotics is ideal when the technology can replace labor-intensive, tedious, or dangerous processes. The ARM Institute, Lockheed Martin Corporation, the University of Southern California, and Texas A&M University have collaborated to develop a robotic sanding and finishing solution to address manual sanding in the production of the Joint Air-to-Surface Standoff Missile. In Troy, AL, Lockheed Martin currently sands 2,000 blades per year and each individual blade takes eight hours. The robotic sanding and finishing approach reduced the time required to sand each blade and increased production rates. Additionally, this new approach saved \$813,000 in the first year of implementation. The ARM Institute assessed that widespread adoption of this technology is expected to yield significant savings and growth over time.

In most circumstances, energetics manufacturing is a labor-intensive, time-consuming process that involves exposure to hazardous chemicals and multiple nonlinear, non-concurrent steps. This long process extends production cycles and makes it logistically challenging to process and finish pieces. Despite the current challenges in energy manufacturing, there has historically been reluctance to change processes and procedures because of concern that new external stimuli could cause energy to be released quickly, disrupting manufacturing and endangering workers and facilities.

ARM worked in conjunction with the U.S. Army DEVCOM AC to develop robotics for handling energetic materials. To reduce the number of touchpoints and potential errors, prime contractor Rensselaer Polytechnic Institute partnered with Ensign Bickford Aerospace & Defense at Picatinny Arsenal in Wharton, NJ, to create an off-the-shelf solution for safely pouring granular energetic materials that reduces both human exposure and potential hazardous dust clouds and touchpoints. This led to a design for a container capable of handling



explosive material, a set of robot code, and a supervisory system controller. This solution removed humans from material handling, increased throughput, and improved end-product quality without requiring production line stoppages.

The team at ARM and U.S. Army DEVCOM AC team also worked with Siemens and University of Southern California to utilize imitation learning, using a robot to accurately mimic operators scooping and weighing fine increments of powder. While performed solely with inert materials, this effort demonstrated an initial capability for incorporating robots into hazardous environments, reducing operators' exposure while performing high concentration and labor-intensive tasks. This effort is being incorporated into safe and automated energetic handling research at U.S. Army DEVCOM AC for future transition to the Industrial Base.

Likewise, ARM is working with Indian Head Naval Surface Warfare Center (NSWC), and the U.S. Army DEVCOM AC to apply robotics and automation technologies to explosive pressing operations. The presses form explosive billets that are used in warhead applications. Integrating automation into this process is aimed to reduce operator exposure to hazardous environments and material handling, as well as improving the overall efficiency of the manufacturing process. Both Indian Head NSWC and the U.S. Army DEVCOM AC are planning to develop and evaluate these technologies in prototype systems. Pending a successful prove out, U.S. Army DEVCOM has found interest in further work to transition production lines at manufacturing sites such as American Ordnance, and NSWC Crane.

Also involving some energetics handling, automation is being applied to produce 155mm munitions at three different facilities. At a General Dynamics-Ordnance and Tactical Systems (GD-OTS) facility in Marion, IL, they are producing bag charges for 155mm weapons systems. The assembly process is currently integrating three robotic stations, and the bags are tied using two articulated robots. Additionally, propellant, black powder, and Clean Burning Igniter are delivered to the production line using an Autonomous Mobile Robot. An articulated robot is responsible for stenciling cans and loading pallets. GD-OTS in Hampton, AR, is also using automation in its M795 line for 155mm production. Armtec in Coachella, California, is deploying robots to automate its production line, improving safety and quality. These lines, now automated, have reduced the labor required for 155mm production from 60 to 80 people to 4 to 8 people.

### **Digital Engineering and Sensing**

DoD Instruction 5000.97 defines digital engineering as the use and integration of digital models and underlying data to support the development, testing, evaluation, and sustainment of a system. Digital engineering can be used to modernize how the Department designs, develops, delivers, operates, and sustains systems. Sensors can complement digital engineering by providing real-time data from factory-floor equipment. Even simple sensors can be deployed to effectively monitor machines and items through sensing. Digital tools can analyze data from sensors to monitor quality and drive real-time process improvement. in real-time.

For example, MxD created a sensor kit that is machine-agnostic and requires no coding for implementation. When deployed, the sensing kit can collect data and monitor manufacturing processes, reducing downtime and failures, including on legacy equipment. Data analysis from the sensor kit enables real-time equipment data to inform and improve operations, enhancing efficiency and productivity at the manufacturing site.

Digital tools and sensing have been particularly successful at the U.S. Army DEVCOM AC, which ran a pilot-scale plant where organic material is processed using a melt-pour process. MxD led an effort to implement sensor kits discussed above with the help of sensor integration experts affiliated with the company Concept Reply. They also chose to partner with XMPRO to provide the U.S. Army DEVCOM AC with a Digital Twin Framework, as well as the



implementation of a digital twin on the existing melt-pour pilot scale production cell at the U.S. Army DEVCOM AC. Before deploying sensors, the melt-pour process at U.S. Army DEVCOM AC relied on an operator watching the operation and using temperature thresholds to determine whether the slurry mixture was ready to be poured. The implementation of the digital twin enabled the operator to make that decision based on data from industrial sensors rather than relying on and manual recordings, which enhanced precision and consistency across pours. This low-cost digital twin was able to derisk and demonstrate successful implementation, which could be scaled across U.S. Army DEVCOM AC and the DIB. MxD is continuing to work with U.S. Army DEVCOM AC to pilot and prototype this new method for the melt-pour process applies to transform and pour it into shells with uniform quality. The decision on when to pour the mixture will be data-driven and rely on sensors that measure factor such as, sonic velocity, surface tension, viscosity, density, and temperature.

### **Advanced Materials**

Advanced composites combine different materials to produce superior properties that the individual components do not possess. These materials are distinguished from traditional composites by their reliance on cutting-edge materials and manufacturing techniques. These compositions include reinforcement, which strengthens the targeted material, matrix, which holds the components together, and gradients, which replace fragile boundaries with blended materials. For munitions, plasma arc welding is a cutting-edge manufacturing technique that allows for deeper and narrower welds than traditional techniques.

LIFT has opened an Advanced Metallic Production and Processing (AMPP) Center in Detroit, Michigan, which is responsible for development-scale production of custom alloys and feedstocks for metals to support advanced manufacturing. AMPP can produce advanced materials, including metallics, ceramics, high-performance polymer composites, and graded materials. These advanced materials can support parts production and reduce reliance on materials with limited availability by providing substitutes that can be advanced manufactured. This approach integrates physics-based modeling, rapid prototyping, and iterative validation by combining materials design, simulation, and small-batch production. LIFT goal is to employ this approach to dramatically reduce the typical materials development and qualification life cycle—from 5–10 years to just 2–3 years.

Current LIFT brick-and-mortar efforts include finding a more printable alternative to tungsten. LIFT is now working on coupled materials and predicting the performance of materials for hypersonic munitions. For legacy munitions, traditional welding techniques are a source of band failure in 155mm shells, and plasma arc welding could improve band reliability and increase manufacturing productivity fourfold.

### **Potential Solutions and Reforms and Best Practices**

The following reforms, discussed and summarized in Table 9, represent options that would improve production capacity and surge capability for legacy munitions by addressing challenges to the adoption of advanced manufacturing. Some of these options are best practices that already contributed to the pathfinders discussed above. Including advanced manufacturing in new developmental programs is a more straightforward process, but it has limited applicability to producing legacy munitions. Similarly, implementing advanced manufacturing approaches is easier in new facilities than in established production lines, though upgrading established lines is particularly important for legacy munition production. Those forms of advanced manufacturing may prove advantageous for new systems, but the scope of this paper focuses on advanced manufacturing that would providing a positive return on investment for legacy systems.



## **Producer Financial Incentive Reforms and Best Practices**

**1.1** *Create steady demand signals to avoid boom-and-bust cycles and encourage investment and consistent funding.*

A repeated problem for the legacy munitions enterprise is a lack of steady demand signals. A history of boom-and-bust-cycles means that the DIB does not see spending increases as indicative of sustained demand. This uncertainty causes both the government and industry to become unsure there will be a return on their investment. Steady demand signals can be communicated through mechanisms such as multi-year contracting for legacy munitions. This approach contractually reinforces the predictability of procurement spending to provide greater confidence in making investments. The downside of this approach is reduced flexibility and less funding available for new systems, but as a result attracts more private capital than the same amount of spending would via shorter term contracts.

**1.2** *Invest in research lab efforts to advance understanding of chemicals and materials to the applied level and open new options to address supply chain vulnerabilities.*

A direct path to addressing producer financial incentives is for the Department to further subsidize R&D efforts, derisking, and piloting. The Department has identified advanced manufacturing as a priority by embedding it into its RDT&E budget, as well as by specifically earmarking money to support additive manufacturing. The downside of this approach is that the government bears the cost and strategic choices and oversight are critical for effectiveness. Commercial industry is the proving ground for advanced manufacturing approaches. However, ongoing support for basic research labs like the Office of Naval Research and the Air Force Research Laboratory can aid in addressing challenges specific to legacy munition adoption. This work allows demonstration of the potential of different advanced manufacturing approaches or desired properties for advanced materials. These research labs also are critical in supporting pathfinding for a specific application for the Services. Future opportunities highlighted by experts included research into producing chemicals through biomanufacturing or understanding fundamental properties of materials. Research that can address supply chain vulnerabilities for critical chemicals, materials, or alloys or that could provide the knowledge needed to enable qualification reforms would be especially beneficial for legacy munitions.

**1.3** *Invest to derisk and pilot application of relevant advanced manufacturing approaches and expand test and demonstration spaces to encourage advanced manufacturing adoption in the main production lines for legacy munitions.*

Developing and maturing manufacturing approaches that are relevant to legacy munitions is rarely sufficient to incentivize adoption of modernized DIB production lines. Decisionmakers are unlikely to modify established production lines unless the technology is becoming obsolete or there is a demand for higher quantities of the legacy munitions than the technology can support. However, the challenges discussed in the subsection on *Narrow Time Windows for Economical Adoption* mean that prudent investments can expand the menu of options available when the opportunity for modernization arises. This method can also aid Program Managers, System Program Offices, and Program Executive Officers in moving fast while respecting the appropriate grounds for their risk aversion.

The Department invests RDT&E, Defense Procurement & Acquisition Policy, and Industrial Base Analysis and Sustainment program funding enables derisking that helps technology and manufacturing approaches reach higher levels of technological and manufacturing readiness, which may lead to more buy-in. Such funding in combination with DIB investments was the source of most of the pilot case studies discussed above. Further investments in test and demonstration spaces would lower the barriers to entry for small businesses and commercial firms and help address limiting factors for the qualification system.



Government funders could oversee and promote these efforts by collecting and analyzing data to demonstrate the return on investment of advanced manufacturing approaches, including cost savings, improved quality, and increased production capacity. The vendors are the ones who are ultimately responsible for delivering the product. However, analysis from piloting and derisking efforts could aid Program Managers in judging whether advanced manufacturing could address an urgent challenge.

***1.4 Encourage vendors and programs to work with relevant MIIIs, Service Research Labs (SRLs), and consortia to facilitate forming partnerships and funding, prototyping, integrating, and certifying advanced manufacturing approaches.***

New vendors and programs can leverage a range of government resources to aid innovation. Similarly, SRLs, development centers, MIIIs, the larger Departmental ManTech enterprise, and related consortia offer a mix of resources and expertise, which can be helpful for validating new approaches, advising on certification, and helping find or connect to contracts. The MIIIs can help new companies connect with or establish partnerships with primes, aiding the integration of new approaches into legacy systems production. The MIIIs also support prototyping and pilot plant facilities and robotics manufacturing hubs, and SRLs can offer access to their facilities through testing service agreements, which lower the barriers to the adoption of new manufacturing processes. Finally, these organizations can aid would-be partners in finding contracts or funding sources, or in addressing challenges such as funding gaps.

***1.5 Use situationally appropriate mechanisms such as OTA and distributing orders directly or via employing consortia or public- private partnerships to enhance incentives for vendors to prototype and implement advanced manufacturing approaches or to ease engineering changes.***

The Department can also enter various contracts and agreements with vendors to incentivize accelerated adoption of advanced manufacturing approaches. When a manufacturing approach offers a positive return on investment, the traditional contracting approach of such as multi-year procurement discussed in recommendation 1.1 can provide companies with sufficient time to profit from adoption. A wider range of incentives could be incorporated through OTAs, which are legally binding agreements that use statutory authorities to allow federal agencies to enter business arrangements that do not have to comply with many of the rules imposed on acquisition contracts. The Department is authorized to employ OTAs for prototyping and production when a vendor has invested their own funds or when there is significant participation by a nontraditional defense contractor that has not adopted the demanding federal acquisition cost accounting standards. These approaches could work indirectly by setting incentives for cheaper, faster performance, or directly by addressing non-recurring engineering costs associated with automation and digitization. These agreements could also include greater flexibility to avoid or mitigate the need for engineering change requests prompted by the adoption of advanced manufacturing approaches.

When the Department is directly contracting for materials, parts, or components in support of legacy munitions, breaking up purchases into smaller order quantities and leveraging consortia could reward new entrants and small businesses that adopt advanced manufacturing. Small order quantities could employ multi-award contracts, although OTAs often work with consortia that are particularly well suited to collaboration and distributed production. Public-private partnerships and consortia can bring multiple vendors together, bid on behalf of multiple vendors, and then distribute the workload. This has the potential to provide a return on investment for various vendors, even if they are not yet able to scale to typical Departmental order quantities.





However, it is worth noting that OTAs do not require inclusion of the Defense Federal Acquisition Regulation Supplement (DFARS) clauses for explosives, and therefore safety standards are potentially at risk. The buying activity or new safety administration/oversight organizations could mitigate that risk by notifying the Defense Contract Management Agency Safety Center, Safety, and Occupational Health Division and making sure the contractor is in compliance with their OTA. This suggestion coincides with the advice in recommendation 3.1 to engage the DoD Explosive Safety Board early in the process. In addition, OTAs and more distributed contracts may require a larger and/or empowered acquisition workforce.

**1.6** *Create or reinforce policies to address complexities associated with data sharing and IP. Create a shareable IP where appropriate, embrace modular open systems approach (MOSA) or containerized munition models, and encourage vendors to address concerns at the outset.*

DFARS rules on IP and data rights ownership can be arcane and, when combined with traditional industrial approaches, can hamper collaboration and licensing. Vendors are often reluctant to share information they view as proprietary. Some MILs, development centers, and depots have sought to develop IP, such as government reference designs for inert solid rocket motor components, that can be used by the OIB, relevant vendors, or all consortium members. For existing IP of special interest to advanced manufacturing, the Department could better incentivize industry to license and share data on material properties. For example, the producer of additive manufacturing equipment may be more interested in winning defense production contracts themselves than sharing data that would enable their customers to potentially compete with them. This leads to this data often remaining in the head of the technician, being only available on paper, or lost entirely. Similarly, as mentioned above, partnership agreements would help multiple vendors come together to fulfill large contract solicitation orders and the quantities requested. However, explicit agreement at the outset on who owns the IP and controls the data is necessary to address concerns about data rights, data sharing, and IP. These issues need to be resolved before entering into partnership agreements and submitting joint bids. Under traditional models, managing IP and data rights often involves legal obligations and difficult negotiations. However, innovative approaches such as containerized munitions or a MOSA can provide templates that ease the sharing of interface information while allowing vendors to safeguard their most important IP, enabling distributed production facilities. Public-private partnerships can foster collaboration, encourage the sharing of resources, and enable the leveraging of capabilities for dual-use purposes.

## **Departmental Process and Regulation Reforms and Best Practices**

**2.1** *Share testing data between Services and avoid requiring duplicative testing to achieve qualification to speed innovation and reduce costs.*

Reforms that can address the lengthy, time-consuming, and fragmented qualification process are central to advancing the use of advanced manufacturing approaches. Manufacturers are disincentivized from working with the Department because they do not want to allocate significant resources to navigating the qualification and certification process, especially if they are not confident that they will receive a favorable outcome. However, the qualification process plays a vital role not only in ensuring weapon system reliability but also in addressing safety risks for manufacturing workers and weapon system operators. This is especially true for energetic material used in propulsion and warheads, which are rightly subjected to strict safety considerations. Advanced manufacturing offers the potential to improve quality assurance and reduce the role of touch labor in dangerous, dirty, and dull manufacturing processes. However, especially for the newest techniques, the imperative of safety and thoughtful risk-acceptance decisions should be at the forefront of any reforms. Given these difficulties, overcoming fragmentation in the qualification process is necessary to avoid duplicate testing and address data gaps.



***2.2 Where safe and practicable, change the qualification process to allow for end-items that meet well-defined standards to be qualified while reducing or removing importance of assessing the manufacturing process used.***

Because of this difficulty, qualification reforms are easiest to pursue by first focusing on inert components or other situations where an inherently safe environment is more readily achievable. At present, the adoption of advanced manufacturing is slowed by the qualification process that places significant weight on qualifying the manufacturing approach. It would be easier to adopt new and iterate advanced manufacturing approaches if the process increased the focus on ensuring that the end item that meets required specifications and performance requirements.

There are multiple ways to enable a greater focus on the end product rather than the process. In some cases, the commercial industry already employing advanced manufacturing approaches, such as the food and beverage industry, can provide confidence when judging if a manufacturing process meets the criteria for being inherently safe. Inert components, such as motor cases and nozzles used in solid rocket motors, are a good starting point. A related concern is that materials may contain impurities or undergo chemical changes during production using an advanced manufacturing approach. Clearly defining and developing criteria for evaluating potential changes in material properties would reduce qualification uncertainty. Where assessment capacity and safety considerations allow, focusing on the output of the part or component rather than qualifying the process would remove a significant obstacle for the adoption of advanced manufacturing.

***2.3 Encourage and accept more data from manufacturers and tooling suppliers to inform the specification and qualification process.***

A key challenge for the qualification process is that the Department lacks key data on advanced manufacturing production equipment and materials. The Department could invest time and money in generating this data for itself. However, accepting more commercial data to guide qualification and specification decisions would encourage manufacturers to collect and provide more data. Commercial data is especially important for additive manufacturing, where additional data could enable broader standards rather than qualifications specific to a particular machine-feedstock combination. The Department could also relax its standards for the data it accepts, allowing industry testing and certification data to be included in the qualification process. Accepting commercial data would save time and effort, though as trust builds, it may be easiest to use this data as a multiplier for Department-collected data rather than a replacement.

***2.4 Support digital approaches and necessary infrastructure investments by accepting sensor data and digital twins.***

The Department could modernize its policies to reinforce the adoption of digital engineering and model-based system engineering. Shifting inspection regimes to rely on sensors, such as laser phallometers with automated defect detection, can enable the adoption of digital twins. As discussed in the section on *Limits of Incremental Adoption*, manufacturers will typically need to invest in network infrastructure to achieve the benefits. As a result, favorably considering the availability of digital twins and sensor data is an important step to reward the adoption of digital manufacturing approaches. Accepting digital twins and favorably considering sensor-based assessments are important steps toward encouraging the adoption of digital engineering approaches. These benefits may be especially important in supporting the development of advanced materials. Integrated tools and materials enable the simulation of the additive manufacturing and the properties of known alloys, which could allow for digital prototypes.



***2.5 Use innovative arrangements to support rapid prototyping, partnerships, and distributed manufacturing, which can lead to faster fulfillment of government needs.***

As discussed in recommendation 1.5, OTAs provide a legal mechanism for the government to create and promote new technologies in partnership with nontraditional defense contractors and commercial firms by reducing barriers to entry and burdensome requirements. OTAs are particularly relevant for advanced manufacturing because many vendors creating novel technologies using advanced manufacturing approaches are not typically federal government or defense contractors and have commercial products. OTAs help facilitate rapid government development and the acquisition of novel solutions, including advanced manufacturing approaches. Besides, OTAs, there are also Broad Agency Agreements, Small Business Innovation Research (SBIR) funds, Cooperative Research and Development Agreements, and Indefinite Delivery, Indefinite Quantity contracts that offer opportunities to promote and leverage new technologies. These agreements and authorities help to transition commercial solutions to the defense marketplace.

**Culture and Workforce Constraint Reforms and Best Practices**

***3.1 Educate communities, particularly the manufacturing workforce, regarding advances in advanced manufacturing relevant to legacy munitions.***

Education about advanced manufacturing can help reduce skepticism about change. User communities that are resistant to change, such as the energetics community, could be introduced to increasingly mature concepts and have their tolerance for newness tested. Bringing program office staff and leadership on board is especially important as they control production dollars. Educating those involved in existing manufacturing processes about the new processes should ease their concerns or identify issues that still need to be addressed. Additional funding for the qualification process to develop safety protocols and risk assessments can reduce reluctance. Advanced manufacturing subject matter experts should evaluate approaches for effectiveness and applicability to specific programs and provide workforce training tailored to each program's context. In addition to education, the Department can pursue this goal by fostering a culture of innovation and experimentation in government and industry, establishing or leveraging Joint Working Groups, supporting for industry conferences and engagements, and work with academia to close skills gaps.

***3.2 Employ data driven decision making to target adoption of advanced manufacturing and overcome skepticism where the benefits are greatest.***

Decisions about whether to implement a new manufacturing approach should be based on whether there is a return on investment in the approach and on that determination. However, the Department lacks the information needed to make this determination most efficiently. Recommendations 2.2 and 2.3 would help to evaluate whether a proposed or piloted approach would improve quality and production capacity and/or reduce costs. If the Department had relevant data available, it would be easier to persuade Program Managers to implement a particular approach.

***3.3 Provide training and workforce development focused on advanced manufacturing to those with manufacturing careers or those interested in manufacturing careers, leveraging MII programs and resources.***

Developing and employing next-generation workforce training programs can expand the available workforce and increase comfort with advanced manufacturing. Workforce training programs should be used to train individuals who are interested in advanced manufacturing. Additionally, those in the present workforce who would be involved with implementing new advanced manufacturing approaches should be provided with specific training. Next-generation



training programs allow the current workforce to maintain their jobs and repurpose labor into roles such as robot maintenance, programming, and control room positions that help sustain advanced manufacturing technologies. To assist with these issues, more manufacturers could partner and work with MIIIs, which have workforce development programs, including training and certificate programs focused on enhancing proficiency in advanced manufacturing. The MIIIs also help grow the workforce by promoting manufacturing jobs, providing opportunities for regional students through partnerships with local community colleges and universities, designing and executing education programs, offering innovative apprenticeships, and providing access to job databases. Additionally, the Office of the Secretary of Defense's Manufacturing Technology Program manages the Manufacturing Workforce Development Program (M-EWD), which advises the MIIIs on talent development. By developing talent, the MIIIs can help address workforce concerns that their jobs will become obsolete by equipping manufacturing workers with the skills needed to implement advanced manufacturing.

## Next Steps

JPAC will apply lessons from this project to foster proactive failure analysis detection and mitigation. JPAC will leverage lessons learned from artificial intelligence to predict and address common production failure modes to better assess failures/defects during processing before they are discovered in the finished product during acceptance testing—saving time and money during production. In addition, this report suggests further opportunities for piloting or scaling the case studies above to better integrate advanced manufacturing into the defense industrial base in hopes of helping to build surge capability and produce legacy munitions faster. These potential pathways focus on automation technologies more broadly, digital engineering and sensing, reforming qualification for additive manufacturing, and flexible manufacturing.

One relevant area for scaling would be the MxD sensor kits, which have been installed at Portsmouth Naval Shipyard. This simple no-code solution could be integrated into legacy munitions production lines. The current focus of the sensor kits is to provide real-time performance tracking and closed-loop real-time data analytics to allow for proactive analysis detection and mitigation. There are many factories, ammunition plants, and systems that would benefit from using sensing to gather data on real-time progress towards an objective. Turning to automation, the paper cites many instances in which advanced manufacturing approaches

Advanced manufacturing has the potential to improve the speed, efficiency, and agility of munitions manufacturing, which could make it easier to surge production. Implementing new approaches may lead to challenges aligning new opportunities with current Departmental processes, paying large upfront costs, and will thus not always be practical for legacy munitions production lines. However, targeted adoption is likely to lead to long-term gains from process improvement, especially when focused on addressing manufacturing challenges and the maturing of defense applications of commercially proven manufacturing approaches. Advanced manufacturing presents a potential solution to an urgent need to surge munitions and strengthen the industrial base, and as such should be viewed as a critical investment in the supply chain. Present efforts to ramp legacy munition production are an opportunity for the Department to leverage partners, MIIIs, and innovators to mitigate the challenges with integrating advanced manufacturing and modernizing technology. Success would allow for advanced manufacturing to become a core component of the future munitions industrial base.



# **Generating Adaptive Capability from Constraints: Aligning Operational Dependency, Requirements, and Competitive Procurement to Sustain Advantage in Maritime Expeditionary Operations**

**Lieutenant Commander Adam Bradley Pierce**—is a USN Explosive Ordnance Disposal Officer, currently stationed at the Navy Expeditionary Warfighting Development Center, in Virginia Beach, VA. He is a graduate of the Naval Postgraduate School (DDM-Financial Management, 2022).  
[adam.b.pierce.mil@us.navy.mil]

## **Abstract**

This research examines how infrastructure-driven operational dependency in maritime expeditionary environments can be translated into scalable capability through alignment across dependency identification, requirements development, and competitive procurement. Contemporary operations are constrained by contested logistics, distributed infrastructure networks, and interdependent systems whose performance cannot be assumed, yet existing processes do not effectively convert dependency-driven risk into actionable requirements or generate responsive industrial output at scale.

To address this gap, the research develops an integrated framework linking geostrategic logistics evaluation and nodal selection (GEO-LENS), functionally defined capability requirements, and variable portfolio contracting (VPC). GEO-LENS identifies and prioritizes dependencies using a structured Operational Dependency Index, generating distributed demand signals that are executed through a competitive, multi-vendor procurement model.

Results demonstrate that dependency is concentrated in high-leverage nodes and that energy functions as a system-level pacing constraint. The VPC model improves performance while reducing cost and supply risk by shifting innovation incentives from government-directed processes to continuous competition among firms. These findings demonstrate that constraint, when identified and translated into demand, becomes a driver of adaptability, innovation, and scale—transforming operational vulnerability from a source of risk into a source of advantage.

## **Background and Problem Statement**

### **The mission was assessed as “low risk”:**

Within 72 hours, a forward maritime expeditionary force was directed to activate a distributed sustainment network across Europe’s high north and eastern flank. Prior planning assumed access to critical infrastructure, including ports in Northern Norway, rail corridors through Sweden and Finland, and logistics hubs extending into the Baltic region. On paper, capacity existed. Contracts were active. Routes were mapped.

### **Execution revealed reality:**

A regional power disruption reduced throughput at a key intermodal hub. Commercial fuel delivery, which was contracted but not prioritized, lagged behind demand. A port remained open but could not sustain offload rates due to power constraints on material handling equipment. Rail movement slowed as competing civilian demand surged. Digital coordination systems remained operational but degraded, forcing manual synchronization across nodes that had never been tested under simultaneous load.

### **The network did not fail—it revealed constraint.**

No single failure stopped operations. Every system functioned—just below the level required. As delays compounded, sustainment timelines slipped beyond operational tolerance.





Commanders were forced to prioritize movement over persistence, trading true sustainment for access.

## Background

Strategic advantage in the current security environment is no longer defined solely by the scale of existing capability, but by how accurately the operating environment is understood and translated into capability. Conventional strategy formation emphasizes adversary capability and intent, aligning objectives and resources against external threats while assuming that infrastructure, logistics, and industrial systems will perform as expected. This assumption is fundamentally flawed. In distributed and contested environments, operational outcomes are governed not only by adversary action, but by the performance of interdependent systems that introduce constraint independent of enemy interference. These dependencies shape feasibility, tempo, and persistence, yet remain insufficiently identified and prioritized in planning. Expanding analysis to systematically incorporate operational dependency reframes these constraints as inherent features of the environment, requiring that they be translated into capability before they manifest as operational failure.

The 2025 National Security Strategy (NSS) underscores the need to “evaluate, sort, and prioritize,” warning that diffuse prioritization undermines long-term effectiveness. Within this framework, the NSS identifies resilient infrastructure, secure supply chains, and a revitalized industrial and energy base as foundational to sustaining national power, emphasizing that “American national power depends on a strong industrial sector capable of meeting both peacetime and wartime production demands” and calling for “the world’s most robust, productive, and innovative energy sector.” (Trump, 2025) At the same time, the NSS highlights the strategic consequences of dependency, noting how shifts in global supply chains and energy access have increased reliance on external actors and intensified competition.

The 2026 National Defense Strategy (NDS) translates these priorities into a defense posture defined by simultaneous, multi-theater competition, contested access to key terrain, and the central role of industrial capacity in sustaining operations. This posture prioritizes the ability to defend the homeland, deter aggression, and regenerate combat power at scale, supported by efforts to “reinvest in U.S. defense production, build out capacity, empower innovators, and adopt new advances in technology” (Hegseth, 2026). Together, the NSS and NDS establish a consistent strategic imperative: the ability to generate, project, and regenerate combat power under constraint (through aligned operational, industrial, and resource systems) is central to sustaining advantage.

At the operational level, maritime expeditionary forces are inherently dependent on distributed infrastructure networks to generate, project, and sustain combat power. The Chief of Naval Operations’ Fighting Instructions emphasize that naval advantage is produced through the integration of “people, infrastructure, and materiel” within the Foundry—the system that “builds, generates, sustains, and modernizes naval power . . . at the speed and scale necessary to win.” This construct reinforces that infrastructure is not a supporting function, but a central component of warfighting effectiveness. Forward naval forces preserve sea lines of communication, enable deterrence, and underpin the economic and logistical flows that sustain national power, operating with “expeditionary reach, unmatched mobility, and persistent presence” in contested environments. In conflict, this dependence intensifies, as the Fleet must “maneuver, strike, and sustain operations across all domains” while absorbing disruption to the very infrastructure upon which it relies (Caudle, 2026).

Similarly, the Commandant of the Marine Corps’ Planning Guidance reinforces the requirement for forces to remain forward, distributed, and expeditionary while accounting for the realities of protracted, high-intensity conflict. The guidance emphasizes that “service planning



accounts for the significant risk of protraction in a peer versus peer conflict,” requiring “sufficient depth of magazine supported by a resilient and distributed logistics network to persist throughout a protracted high-intensity fight.” At the same time, it recognizes that the character of war is evolving through extended-range kinetic capabilities and increasingly sophisticated non-kinetic effects across space, cyberspace, and the electromagnetic spectrum, demanding that forces “be increasingly creative” in how they train and operate (Smith, 2024). Together, these priorities reinforce that operational effectiveness depends not only on forward presence, but on the ability to sustain distributed forces through resilient infrastructure and adaptive support systems in contested environments.

These concepts converge on a common operational reality: maritime expeditionary operations are only as effective as the infrastructure networks that enable them. This dependency extends beyond traditional logistics to include energy generation and distribution, maintenance and repair capacity, transportation nodes and access corridors, and the digital and electromagnetic systems required for command and control.

The 2024 Regional Sustainment Framework (RSF) identifies the need to transition from centralized sustainment toward distributed, regionally integrated networks that leverage allies, partners, and industry (LaPlante et al., 2024). While this approach enhances survivability and operational reach, it introduces a critical tension: distribution increases resilience while simultaneously increasing dependence on a wider and more fragile network of infrastructure and supply chains. Operational effectiveness therefore becomes increasingly sensitive to disruptions in access, throughput, sustainment capacity, and enabling resources such as energy.

The Navy Expeditionary Warfighting Development Center’s Advanced Naval Basing (ANB) Innovation Forum provides empirical grounding for this tension by identifying early-phase establishment as the period of greatest vulnerability, when assumptions compound and resilience is lowest (Reitter, 2026). The Forum highlights that infrastructure constraints, spread across access, throughput, sustainment, and contractor dependence, do not emerge in isolation but interact to degrade operational feasibility during initial force projection, reinforcing a defining condition:

**Constraint, not capacity, defines the operational environment.**

Compounding these challenges are structural limitations within defense acquisition and procurement systems. Portfolio Acquisition Executive (PAE) transition memoranda and related reform efforts highlight persistent issues, including slow capability delivery timelines, misalignment between requirements and acquisition outcomes, fragmented decision-making structures, and insufficient adaptability to emerging operational needs (Potter et al., 2025). Although recent reforms emphasize iterative capability development, modular open systems approaches, and improved alignment between requirements and acquisition, the system remains constrained in its ability to translate operational need into scalable industrial response.

**Problem Statement**

Across strategic guidance, operational doctrine, sustainment frameworks, and acquisition reform efforts, a consistent pattern emerges: operational dependency (defined as the measurable reliance of military operations on the performance and resilience of interconnected infrastructure, logistics networks, and enabling systems, which collectively determine the ability to generate, sustain, and project combat power under constraint) is increasing, yet the ability to translate that dependency into capability remains insufficient. Existing processes do not effectively convert dependency-driven risk into actionable requirements, and procurement systems lack the flexibility and competitive structure necessary to adopt innovations rapidly and generate capability at scale.



Figure 1 depicts the *Spectrum of Operational Dependency*, or the conceptual continuum that defines the range of possible reliance on organic versus external infrastructure, logistics networks, and enabling systems, through which commanders and planners assess and optimize trade-offs between risk, cost, responsiveness, and scalability in order to sustain operational effectiveness under constrained conditions.




**Figure 6. Spectrum of Operational Dependency**

Accordingly, the central problem addressed in this research is that the Department of Defense lacks an integrated framework to identify, prioritize, and respond to infrastructure-driven operational dependencies through adaptive capability development and procurement. Without such a framework, constraints imposed by contested environments remain latent until they manifest as operational failure—existing as unidentified and unmanaged risk.

**This research addresses that gap by asking:**

How can operational dependency, often unmeasured in conventional planning but decisive under contested conditions, be systematically translated into prioritized capability requirements and fulfilled through adaptive, competitive procurement mechanisms?

**Scope and Goals**

This research examines how operational dependency, often unmeasured in conventional planning but decisive in execution, can be translated into scalable capability through aligned requirements and competitive procurement in maritime expeditionary environments.

The analysis focuses on distributed maritime expeditionary operations conducted in contested environments, where forces rely on a combination of host-nation infrastructure, commercial and dual-use systems, and military-owned capabilities. These environments are characterized by contested access, degraded logistics networks, and limited reliance on centralized sustainment, conditions that amplify infrastructure dependency.

The research centers on infrastructure-driven operational dependency rather than platform-centric limitations. This includes transportation networks, access corridors, sustainment capacity, energy systems, contracted services, and digital infrastructure. These dependencies are treated as measurable and comparable variables through the geostrategic logistics evaluation & nodal selection (GEO-LENS) construct.

The research is bounded to high-quantity, potentially high-attribution operational enablers, such as distributed energy systems, unmanned platforms, sensors, munitions, and expeditionary sustainment assets, where demand is scalable and loss is expected, rather than major defense acquisition programs characterized by low production volume and long development cycles. This reflects the need for scalable, modular, and rapidly fielded capabilities aligned with contested environments (Potter et al., 2025).

Methodologically, the research evaluates an integrated framework consisting of dependency identification (GEO-LENS), requirements translation, and competitive procurement (variable portfolio contracting; VPC).

The research pursues three goals:

1. Identify and prioritize operational dependencies
2. Translate those dependencies into functional requirements
3. Evaluate procurement structures capable of delivering scalable capability under constraint.

## **Summary of Literature and Foundational Constructs**

A growing body of research across defense analysis, infrastructure systems, and industrial strategy converges on a central insight: modern conflict is increasingly defined by contested logistics and infrastructure dependency.

Contemporary operational analysis underscores that logistics has shifted from a supporting function to the decisive factor in great power competition. As Zachary Hughes argues, in conflicts between peer adversaries, “logistics is likely to be the primary determinant of military success or failure,” reframing warfare itself as a contest of “competitive logistics.” While distributed operations enhance survivability, they introduce significant sustainment complexity, particularly when infrastructure is degraded or actively contested. Efforts to mitigate vulnerability, such as dispersion, prepositioning, and push-based resupply impose what Hughes describes as “compounding inefficiencies,” where disruption at any point in the supply chain propagates across the system and degrades overall effectiveness (Hughes, 2024). Under these conditions, logistics is no longer a problem of throughput optimization, but a system-level constraint that governs the feasibility, tempo, and persistence of operations across interconnected nodes.

Infrastructure must therefore be understood not as a collection of independent assets, but as an integrated system whose value emerges from networked interaction. The European Union’s Trans-European Transport Network (TEN-T) provides a practical model for this perspective, emphasizing multimodal integration, interoperability, and the elimination of bottlenecks across a distributed transport architecture. As codified in the 2024 TEN-T regulation, transport infrastructure “functions as a network,” where even the “non-operability of a small segment can hamper the efficiency and competitiveness of the system as a whole” (Publications Office of the European Union [POEU], 2024). This framing shifts the focus from individual node capacity to the performance of connections and flows across the network. Accordingly, system effectiveness is determined not by the strength of isolated assets, but by the resilience and coherence of their interactions. This systems-based understanding aligns directly with the GEO-





LENS construct, which evaluates infrastructure as a network of interdependent capabilities whose collective performance defines operational feasibility.

The NEXWDC Advanced Naval Basing Innovation Forum provides empirical validation of these theoretical constructs. The Forum demonstrated that distributed basing concepts are most vulnerable during early establishment phases, when infrastructure is incomplete and resilience is lowest. It identified multiple interdependent constraints, including access limitations, throughput restrictions, sustainment capacity gaps, contractor reliance, and digital infrastructure vulnerabilities. These constraints do not operate independently; they interact in ways that compound risk and reduce operational effectiveness. The Forum further highlights that energy and energy control systems function as system-enabling dependencies, shaping the performance of all other operational functions (Reitter, 2026).

The RSF reinforces a fundamental shift from centralized sustainment toward a “globally connected, resilient . . . ecosystem” built on distributed, collaborative networks that integrate allies, partners, and the defense industrial base. Rather than relying on traditional models that retrograde materiel to centralized locations, the framework emphasizes positioning maintenance, repair, and sustainment capabilities “closer to the point of need” within a network of geographically dispersed nodes (LaPlante et al., 2024). This transition enhances survivability, responsiveness, and operational flexibility in contested logistics environments, but it also expands the number of critical dependencies across infrastructure, industrial capacity, and partner-enabled supply chains. As sustainment becomes distributed, the system’s performance is increasingly determined by the interaction of these nodes rather than any single capability. Consequently, operational effectiveness becomes more sensitive to disruption, reinforcing the requirement for analytical approaches that can identify, prioritize, and manage dependency across interconnected networks—an imperative directly addressed by the GEO-LENS construct.

At the national level, analysis of the defense industrial base reveals a structural misalignment between operational demand and industrial capacity. Contemporary supply chains remain optimized for efficiency and cost minimization rather than surge production and sustained wartime output, limiting their ability to respond to rapid increases in demand. As highlighted in *Revitalizing National Security Supply Chains*, the “DoD faces a disconnect between its operational planning and its industrial planning,” a gap that constrains the nation’s ability to translate requirements into deliverable capability at scale (Attar et al., 2025). This misalignment is compounded by insufficient integration across government, industry, and allied partners, as well as limited visibility into real-time demand signals and production capacity. Without a coherent mechanism to generate, communicate, and sustain demand signals, industrial systems cannot effectively align investment, workforce, or production decisions to operational needs. The result is a system in which demand consistently outpaces capacity under contested conditions, reinforcing the requirement for frameworks that can translate operational dependency into actionable requirements and align procurement structures to drive responsive, scalable industrial output.

The Department of Defense’s Office of Strategic Capital (OSC) reinforces the central role of private capital, market competition, and financial incentives in scaling national security capability. Rather than relying solely on government-directed investment, the OSC framework recognizes that “private markets continue to finance most investments” in critical technologies and industrial capacity, making engagement with private firms essential to generating competitive advantage (Austin, 2025). This perspective reframes industrial responsiveness as a function of capital flows, risk-adjusted returns, and competitive positioning across economic networks, key industries, and critical technologies. As detailed in the FY2025 Investment Strategy, strategic competition increasingly unfolds within capital markets themselves, where investment decisions shape long-term industrial capacity and technological leadership.





Consequently, procurement and investment mechanisms must do more than acquire capability; they must actively influence industrial behavior by structuring incentives, shaping competition, and directing capital toward priority outcomes. This shift underscores a critical requirement: the design of procurement systems must align with market dynamics to mobilize private investment, accelerate production scaling, and sustain competitive advantage in contested environments.

Taken together, these bodies of work establish several foundational conditions. Operational effectiveness is increasingly constrained by infrastructure and logistics systems. These constraints are systemic and interdependent, often propagating across networks rather than remaining localized. Existing processes do not consistently translate dependency-driven risk into actionable requirements, and procurement systems are not structured to respond to those requirements at the speed and scale required.

### **These findings define the foundational requirement for this research:**

A framework is required to identify operational dependencies, translate them into capability requirements, and align procurement mechanisms to deliver scalable capability under constraint.

## **Conceptual Framework**

This research proposes a framework that transforms operational dependency into capability through alignment across three functions: dependency identification, capability requirements, and competitive procurement. Rather than treating capability development as a linear acquisition process, the framework defines it as a system in which operational conditions generate demand, and aligned mechanisms produce capability in response.

The GEO-LENS construct, provides the dependency identification function. GEO-LENS evaluates geographic nodes and infrastructure systems using the Operational Dependency Index (ODI), which incorporates factors such as infrastructure capacity, access, utilities, supply chains, governance, digital systems, environmental conditions, and threat exposure. These factors are normalized and combined to produce comparative assessments of nodes within a distributed network. The output is a prioritized set of dependencies expressed in terms of node-level risk, contribution to operational objectives, and interdependence with other nodes.

Within this framework, GEO-LENS defines where and why capability is required. It does not prescribe solutions; instead, it establishes the operational conditions that generate demand. This distinction is critical, as it ensures that capability development is driven by dependency rather than pre-determined solutions.

Capability requirements serve as the translation mechanism between dependency identification and procurement execution. Requirements convert identified dependencies into functional descriptions of what capabilities must achieve to mitigate operational risk. In this framework, requirements are expressed in terms of operational effects (such as enabling throughput, sustaining operations, maintaining survivability, or supporting command and control) rather than specific technical configurations. This approach enables modularity, scalability, and adaptability, allowing multiple technical solutions to satisfy operational demand.

The structure of requirements directly influences the effectiveness of procurement. Functionally defined requirements create clear demand signals that can be addressed by a range of industrial providers, while overly prescriptive requirements constrain competition and limit adaptability. As such, requirements function as the central mechanism through which operational dependency is translated into actionable industrial demand.

Competitive procurement, implemented through variable portfolio contracting, provides the execution function. The VPC model engages multiple vendors simultaneously and allocates



production dynamically based on cost, schedule, and performance. This structure maintains continuous competition, incentivizes improvement, and distributes production across multiple providers. By avoiding reliance on a single supplier, it reduces supply risk and enables surge capacity.

Within the framework, procurement is not a passive mechanism for acquiring capability but an active tool for shaping industrial behavior. The ability to adjust allocation based on performance creates a feedback loop in which industrial output evolves in response to operational demand signals and competition between contracted firms.

The interaction of these three elements forms a capability pipeline (depicted in Figure 2).




**Figure 7. Aligned Capability Pipeline**

Alignment across this system determines outcomes. When dependency identification accurately reflects operational conditions, requirements translate those conditions into clear demand signals, and procurement enables responsive industrial execution, the system produces scalable and adaptive capability. When misaligned, the system produces delays and inefficiencies which translate to increased operational risk.

This framework establishes that capability is not assumed. It is produced through the alignment of systems operating under constraint.

The interaction of dependency identification, requirements development, and procurement execution establishes a system in which operational conditions generate demand and aligned mechanisms produce capability. Within this system, the role of each component is distinct but interdependent.

GEO-LENS provides a structured method for identifying and prioritizing dependencies across a distributed infrastructure network. By evaluating nodes against operational objectives and constraint factors, it produces a prioritized set of dependencies that define where capability is required and the relative importance of those requirements. Importantly, this output reflects both individual node performance and the interdependence of nodes within the broader system.

Requirements translate these prioritized dependencies into functional demand signals. This translation is not a direct mapping of infrastructure limitations to technical solutions; rather, it defines the operational effects required to mitigate risk. By focusing on function rather than form, requirements enable multiple pathways to capability development and support adaptability as conditions evolve (Pierce, 2025).

Procurement executes these requirements through competitive mechanisms designed to align industrial behavior with operational demand. Within the VPC construct, multiple vendors compete continuously across cost, schedule, and performance dimensions. Allocation of production is adjusted dynamically, ensuring that industrial output reflects both demand signals

and contractor performance. This structure reduces reliance on single providers, increases resilience, and enables scalable production (Pierce, 2025).

A critical implication of the VPC model is the shift in innovation incentives from the government to the industrial base. Under traditional contracting structures, innovation is largely government-directed, requiring predefined requirements, dedicated funding, and extended development timelines. In contrast, the VPC structure creates continuous post-award competition in which firms gain or lose production share based on cost, schedule, and performance outcomes. This dynamic incentivizes firms to independently invest in process improvements, production efficiency, and technological innovation in order to increase competitiveness. As a result, innovation becomes decentralized, self-funded, and continuous, rather than episodic and government-driven. This shift not only accelerates the incorporation of innovation into operational capability, but also aligns industrial investment with operational demand signals, reinforcing the adaptability and scalability of the procurement system.

The system is inherently iterative. Changes in operational conditions alter dependency assessments, which in turn adjust requirements and influence procurement decisions. Similarly, industrial performance feeds back into procurement allocation, shaping future production capacity and capability outcomes. This feedback loop ensures that capability development remains responsive to changing conditions rather than fixed at a single point in time.

The framework therefore establishes capability generation as a dynamic process driven by alignment:

Dependency identification defines demand; requirements structure that demand; procurement delivers capability at scale.

## **Methodology**

This research employs a structured, two-phase experimental design to evaluate the effectiveness of the proposed framework. The design separates dependency identification from procurement execution, allowing independent assessment of decision-making processes and industrial response. This separation enables analysis of both individual components and the effects of alignment across the system.

## **Operational Scenario and Analytical Context**

### **Area of Operations (AO)**

The experiment is conducted within a contested maritime expeditionary environment spanning Northern Europe (incorporating Arctic logistics corridors, Baltic and Northern European nodes) and Black Sea and Eastern Mediterranean access points, capturing a region characterized by dense infrastructure networks, strategic chokepoints, and exposure to real-world geopolitical constraints.

This AO is selected because it:

- Contains interdependent infrastructure systems (ports, rail, energy, logistics nodes)
- Reflects operationally relevant geopolitical friction
- Introduces scarcity across access, throughput, and sustainment functions

Within this environment, infrastructure operates simultaneously as an enabler and a constraint. Operational effectiveness is therefore dependent on network performance rather than isolated node capacity.



## Strategic and Threat Context

The scenario incorporates a competitive environment defined by a near-peer adversary, the Russian Federation, whose strategic posture introduces friction across the same infrastructure systems required for maritime expeditionary operations.

Adversary effects are represented across four primary vectors:

- Infrastructure contestation: disruption or degradation of ports, rail, energy, and repair nodes
- Arctic leverage: control of northern maritime routes and energy access
- Black Sea access denial: constraints on maritime throughput and sustainment
- Industrial endurance: sustained operations under attrition and sanction conditions

This context aligns with the National Defense Strategy's identification of Russia as a persistent threat across the Eastern Flank, Arctic, and Black Sea regions.

## Strategic Objectives

The experiment is structured around three prioritized strategic objectives, providing the basis for evaluating infrastructure dependency and capability demand across the operational network:

1. Preserve NATO freedom of action on the Eastern Flank
2. Deny adversary leverage over Arctic and High North access
3. Limit adversary ability to sustain protracted conflict through industrial and energy pressure

## Time Horizon

The experiment is conducted over a 36-month period (February 1, 2026–January 31, 2029), a duration enables observation of both short-term friction and longer-term system behavior:

- Initial establishment under constrained conditions
- Evolution of infrastructure dependency over time
- Contractor performance and adaptation
- Procurement reallocation effects

## Phase I: Dependency Identification and Prioritization (GEO-LENS)

Phase I applies the GEO-LENS model to identify and prioritize infrastructure-driven dependencies. A set of operational nodes is selected based on geographic relevance, infrastructure diversity, and contribution to strategic objectives.

A set of geographically distributed nodes is defined to represent key infrastructure within the operational environment. These nodes include ports, airfields, logistics hubs, rail corridors, and regional sustainment and energy systems, including (depicted in Figure 3):

- Central Europe Transit Corridor: Ostrava–Žilina–Budapest (overland logistics corridor)
- Baltic Region: Klaipėda Port; Riga Logistics Hub
- High North / Arctic: Narvik Port; Bodø Air & Multimodal Node; Luleå Intermodal Hub; Oulu Sustainment Center
- Black Sea: Varna Port
- Eastern Mediterranean: Alexandroupoli Port








**Figure 8. GEO-LENS Nodes and Objectives**

Each node is evaluated using the Operational Dependency Index (ODI), which incorporates a standardized set of assessment factors, which are normalized and aggregated to produce node-level dependency scores, enabling comparison across the network:

- Infrastructure capacity and transportation networks
- Utilities and energy availability
- Supply chains and contracted services
- Governance and regulatory environment
- Digital and command and control (C2) systems
- Environmental and climate factors
- Threat exposure

The ODI inputs (generated from United States Army Corps of Engineers Theater Infrastructure Framework and Country Profile documents and open-source information) are normalized and processed within the GEO-LENS model to produce node-level dependency scores. The model evaluates nodes relative to defined strategic objectives and generates a prioritized set of infrastructure nodes. In addition to prioritization, the model produces a resource-constrained synchronization matrix that identifies the sequencing and coordination of actions required to support operational objectives.

The output of Phase I is a structured representation of operational dependency, expressed in terms of prioritized nodes, identified constraints, and associated risk. This output



defines where capability is required but does not prescribe how that capability should be delivered.

### **Requirements Translation: From Dependency to Demand**

Following dependency identification, GEO-LENS outputs are translated into capability requirements. This process converts identified constraints into functional descriptions of what capabilities must achieve to mitigate operational risk.

Requirements are expressed in terms of operational effects rather than technical specifications. These effects include enabling throughput across constrained transportation networks, sustaining operations at distributed nodes, maintaining survivability in degraded environments, and supporting C2 functions under contested conditions. This approach ensures that requirements remain adaptable and do not constrain solution space unnecessarily.

The translation process produces a demand signal characterized by three key attributes. First, it is distributed, reflecting the geographic dispersion of prioritized nodes. Second, it is scalable, allowing capability quantity and allocation to adjust based on operational conditions. Third, it is modular, enabling multiple technical solutions to satisfy the same functional requirement.

This demand signal provides the basis for procurement execution in Phase II.

### **Phase II: Competitive Procurement and Industrial Response (VPC)**

Phase II applies the VPC model to simulate industrial response to the demand signal generated in Phase I. The VPC model is designed to evaluate how procurement structure influences capability delivery under conditions of competition and uncertainty.

A set of contractors is defined, each characterized by production rate, unit cost, and performance attributes (depicted in Figure 4). These contractors represent heterogeneous industrial actors with varying capabilities and constraints. Production is allocated across contractors simultaneously, rather than through a single-award structure.

#### ***Figure 9. VPC Contractor Rating Calculator***

Allocation decisions are based on cost, schedule, and performance metrics. As contractors perform over time, their relative performance influences future allocation, creating a dynamic system in which market share shifts in response to demonstrated capability. This structure sustains competition and incentivizes continuous improvement.

The model incorporates variability in contractor performance, reflecting changes in production efficiency, supply chain conditions, and operational disruptions. These dynamics enable assessment of how procurement structures respond to uncertainty and how they influence supply resilience and scalability.

The VPC model is implemented within a structured analytical environment that tracks contractor performance, allocation decisions, and capability delivery over time. This allows for comparative analysis across different weighting schemes, including even-weighted, cost-weighted, schedule-weighted, and performance-weighted allocation strategies.



## Integration of Phases

The two phases are integrated through the linkage between dependency identification, requirements translation, and procurement execution. GEO-LENS outputs define where capability is required, requirements translate those needs into demand signals, and VPC execution determines how effectively those signals are fulfilled.

This integration enables evaluation of the central thesis: that Constraint becomes a driver of capability when operational dependencies are translated through aligned requirements and fulfilled through competitive procurement, producing measurable improvements in capability outcomes.

## Evaluation Criteria

The effectiveness of the framework is assessed across four dimensions. Decision quality is evaluated based on the ability of GEO-LENS to identify and prioritize meaningful dependencies. Procurement adaptability is assessed through the responsiveness of the VPC model to changes in contractor performance. Supply resilience is measured by the ability to maintain production under variability and disruption. Capability outcomes are evaluated in terms of performance, cost, and delivery over time.

These criteria provide a comprehensive assessment of how alignment across the framework influences the generation of capability under constraint.

## Results and Analysis

### GEO-LENS Outputs: Dependency Prioritization

Application of the GEO-LENS model produced a prioritized set of infrastructure nodes aligned to defined strategic objectives within the European operational environment. The results demonstrate that operational dependency is not evenly distributed across the network, but instead concentrated within a limited number of high-leverage nodes whose performance directly enables access, throughput, sustainment, and network integration.

Across multiple strategic objectives, GEO-LENS consistently identified Klaipėda Port and the Central Europe Transit Corridor as repeat selections, indicating that these nodes function as theater-critical infrastructure. Their capacity and resilience enable NATO reinforcement of the Eastern Flank while simultaneously constraining adversary sustainment under protracted conflict conditions. This convergence highlights that certain nodes are not only important, but disproportionately decisive within the operational system.

Beyond individual prioritization, the model reveals that infrastructure operates as an interconnected system organized into complementary geographic and functional clusters. Three primary infrastructure groupings emerge:

- Baltic / Eastern Flank Cluster: Klaipėda Port, Riga Logistics Hub, Central Europe Transit Corridor
- High North / Arctic Cluster: Narvik Port, Oulu Sustainment Center, Bodø Multimodal Node
- Southern Access Cluster: Alexandroupoli Port

These clusters demonstrate that operational effectiveness depends on coordinated performance across distributed, multimodal nodes. Degradation or constraint at a single node propagates across the network, reducing overall system effectiveness. As a result, infrastructure must be evaluated collectively rather than as isolated assets.



The GEO-LENS outputs establish two defining conditions. First, operational dependency is concentrated within specific nodes and network clusters rather than distributed uniformly. Second, the value of these nodes is derived not only from individual capacity, but from their role within a system of interdependencies. Operational effectiveness is therefore governed by network performance, reinforcing the requirement to prioritize resilience and capacity across interconnected infrastructure systems rather than optimizing individual nodes in isolation.

### **Constraint Identification Across Infrastructure Systems**

The translation of GEO-LENS outputs into operational analysis revealed a range of infrastructure-driven constraints that shape the feasibility of maritime expeditionary operations. These constraints include transportation throughput limitations, access restrictions, sustainment capacity gaps, contractor and labor dependencies, and vulnerabilities in digital and C2 infrastructure.

Transportation constraints were observed in multimodal corridors and logistics hubs where port offload rates, rail capacity, and transfer efficiency limited the movement of materiel. Access constraints emerged from a combination of geographic, political, and regulatory factors that affected the ability to utilize infrastructure in contested or uncertain environments. Sustainment constraints reflected uneven distribution of maintenance and repair capabilities, particularly at forward nodes where capacity to preserve operational readiness is limited. Contractor dependence introduced variability in labor availability and scalability, while digital infrastructure vulnerabilities affected coordination, information flow, and C2 effectiveness.

These constraints did not manifest uniformly across all nodes. Instead, their impact varied based on geographic location, infrastructure maturity, and operational context. Some nodes were primarily limited by throughput, while others were constrained by access or sustainment capacity. This variation reinforces the need for structured analysis to identify and prioritize constraints rather than relying on generalized assumptions.

### **Constraint Convergence: Energy as the System-Level Driver**

While multiple constraints influenced operational performance, analysis across all prioritized nodes revealed a consistent pattern of convergence around a single dependency. Energy availability emerged as a cross-cutting constraint affecting all operational functions and infrastructure nodes within the system.

Energy underpins transportation throughput by enabling port operations, material handling equipment, and logistics networks. It supports sustainment by powering maintenance, repair, and life-support systems at distributed nodes. It enables C2 functions through digital and communication infrastructure. It is also essential for survivability, particularly in austere or degraded environments where external support is limited.

Unlike other constraints, which varied in significance across nodes, energy consistently influenced the effectiveness of all operational functions. This convergence indicates that energy functions as the system-level pacing constraint governing the performance of the broader infrastructure network. While other constraints shape operational risk in localized or context-specific ways, energy defines the baseline conditions under which all other capabilities operate.

This finding aligns with empirical observations from the Advanced Naval Basing Innovation Forum, which identified energy and energy control systems as foundational to distributed operations. More broadly, it reinforces the principle that constraint convergence, rather than isolated limitation, determines system-level performance.



## Requirements Formation: Translating Constraint into Demand

The identification of convergent constraints enabled the development of a structured demand signal through requirements translation. By expressing requirements in functional terms, the analysis converted infrastructure-driven dependencies into capability needs aligned with operational conditions.

The resulting demand signal exhibited three defining characteristics. First, it was distributed, reflecting the geographic dispersion of prioritized nodes and infrastructure clusters. Second, it was scalable, allowing capability quantity and allocation to adjust based on operational tempo and resource availability. Third, it was modular, enabling multiple technical solutions to satisfy functional requirements without constraining innovation.

These characteristics ensure that requirements capture operational need without prescribing specific technical solutions. This approach expands the solution space available to industry and supports adaptability as conditions evolve. More importantly, it establishes requirements as the mechanism through which dependency is translated into actionable demand.

## Competitive Procurement Outcomes: Variable Portfolio Contracting

Application of the VPC model to the generated demand signal produced consistent improvements across performance, cost, and scalability metrics. Across multiple allocation strategies (cost-weighted, schedule-weighted, performance-weighted, and balanced models) the results demonstrate that competitive, multi-vendor procurement can simultaneously enhance capability performance while reducing cost.

Performance outcomes improved across all models, with capability performance increasing by approximately 3.9% to 4.1% relative to baseline conditions. Concurrently, average unit costs decreased by approximately 0.47% to 1.04%, while average contract costs declined by approximately 1.05% to 1.22% (depicted in Table 1). These results indicate that performance improvement and cost efficiency are not mutually exclusive under competitive procurement conditions, but can be achieved concurrently through structured competition.

*Table 1. VPC Performance Metrics*

Beyond aggregate performance gains, the VPC model reveals distinct patterns in industrial behavior driven by allocation strategy. Under cost- and schedule-weighted models, production tends to concentrate among high-efficiency contractors, yielding lower unit costs (approximately \$208,000–\$209,000) but producing wider contract value distributions (approximately \$8.7 million–\$20.1 million). This concentration improves efficiency but increases reliance on a smaller subset of suppliers. In contrast, performance-weighted allocation produces a narrower contract value range (approximately \$13.4 million–\$15.8 million), sustaining a broader base of active suppliers and reducing concentration risk. Balanced allocation strategies provide a compromise between these outcomes, maintaining competition while preserving cost discipline (all depicted in Table 2).



**Table 2. VPC Contract Comparison**

These results demonstrate that procurement structure directly influences both industrial behavior and system resilience. By maintaining multiple active suppliers and dynamically adjusting allocation based on cost, schedule, and performance, the VPC model reduces supply risk, sustains competition, and enables scalable production capacity. Importantly, minimum contract values across all models remain sufficiently large to incentivize firms to invest in production tooling, workforce development, and capability improvement, reinforcing long-term industrial responsiveness.

Collectively, the VPC outputs confirm that portfolio-based procurement not only improves acquisition outcomes, but also shapes the industrial base in ways that enhance both efficiency and resilience under conditions of sustained demand and uncertainty.

### **Integrated Analysis: Alignment as the Determinant of Outcome**

The combined results of GEO-LENS, requirements translation, and competitive procurement demonstrate that alignment across these elements determines whether constraint produces capability or operational risk. Each component performs a distinct function within the system, but their effectiveness depends on how they interact.

GEO-LENS defines the operational conditions by identifying and prioritizing dependencies. Requirements translate those conditions into structured demand signals that reflect operational need. Procurement executes those signals by shaping industrial response through competitive mechanisms. When aligned, these elements form a coherent system that produces scalable and adaptive capability.

Misalignment at any stage disrupts this process. If dependencies are not accurately identified, requirements fail to capture true operational need. If requirements are poorly structured, demand signals become unclear or overly restrictive. If procurement mechanisms are rigid or concentrated, industrial response cannot scale effectively. In each case, the result is delayed capability delivery and increased operational risk.

The results therefore validate the central thesis of this research.

Constraint becomes a driver of capability when operational dependencies are translated through aligned requirements and fulfilled through competitive procurement, producing measurable improvements in capability outcomes.

This reframes constraint from a limiting condition to a generative input. Under this framework, the presence of constraint does not reduce capability; rather, it defines the conditions under which capability must be produced. The ability to identify, translate, and align responses to those constraints determines operational advantage.

### **Summary of Key Findings**

The results yield several key findings. Infrastructure dependency is concentrated in high-leverage nodes and networks rather than distributed evenly. Multiple constraints shape operational performance, but their impact varies across nodes. Energy availability consistently





emerges as the system-level pacing constraint affecting all operational functions. Requirements serve as the critical bridge between dependency identification and industrial response. Procurement structure shapes industrial behavior, influencing both efficiency and resilience. Finally, alignment across dependency, requirements, and procurement determines whether capability is generated effectively under constraint.

## **Conclusion and Recommendations**

### **Conclusion**

This research examined how infrastructure-driven operational dependency and vulnerability—often unmeasured in conventional planning but decisive in execution—can be translated into scalable capability through the alignment of dependency identification, requirements development, and procurement execution. The findings demonstrate that generating capability under constraint is not a function of existing capacity alone, but of how effectively dependency is identified, prioritized, and converted into actionable demand and responsive industrial output.

First, the application of GEO-LENS confirmed that infrastructure dependency can be systematically identified, measured, and prioritized within a distributed operational environment. The model demonstrated that dependency is concentrated within a limited number of high-leverage nodes and networks, and that operational effectiveness is governed by the performance of these interconnected systems rather than individual infrastructure elements.

Second, the translation of these dependencies into functional capability requirements established a clear mechanism for converting operational conditions into structured demand signals. By defining requirements in terms of operational effects rather than technical specifications, the framework enables adaptability, scalability, and competition, ensuring that capability development remains aligned with evolving conditions.

Third, the application of the VPC demonstrated that procurement structure directly influences capability outcomes. Competitive, multi-vendor procurement reduced supply risk, improved performance, and maintained cost discipline while enabling scalable production. These results confirm that procurement is not a passive mechanism for acquiring capability, but an active tool for shaping industrial behavior and aligning production with operational demand.

Taken together, these findings support the central conclusion of this research.

Adaptive capability is generated through alignment across dependency identification, requirements development, and competitive procurement under conditions where constraint, not capacity, defines the operational environment.

This conclusion reframes constraint from a limiting factor to a defining input. Rather than reducing capability, constraint establishes the conditions that drive demand and shape effective responses. The ability to identify, translate, and align responses to constraint determines whether operational systems produce capability or experience degradation.

### **Limitations**

Several limitations should be considered when interpreting the results of this research:

The experimental design relies on modeled environments, including GEO-LENS dependency assessments and VPC-based procurement simulations. While these models provide structured and repeatable analysis, they do not fully capture the variability and complexity of real-world operational and industrial systems. Contractor behavior, supply chain dynamics, and regulatory constraints may differ in practice.



The VPC model simplifies elements of acquisition, including funding timelines, contract negotiation processes, and regulatory requirements. Although the model is designed to reflect competitive dynamics, its implementation within existing acquisition frameworks would require further validation and adaptation.

The accuracy of GEO-LENS outputs is dependent on the quality and completeness of input data. Variability in data sources or assessment assumptions may influence prioritization outcomes, particularly when applied across different geographic regions or operational contexts.

Additionally, the research focuses on a specific operational environment and a bounded set of capability requirements. While the framework is designed to be broadly applicable, further testing is required to assess its effectiveness across different theaters, mission sets, and capability categories.

Despite these limitations, the methodology provides a consistent and analytically rigorous basis for evaluating the relationship between dependency, requirements, and procurement. It enables controlled assessment of how alignment across these elements influences capability generation under constrained conditions.

### **Recommendations for Future Research, Modeling, and Testing**

To expand the applicability of the framework and address identified limitations, future research should extend GEO-LENS analysis to additional operational environments, particularly within the Indo-Pacific, where geographic dispersion, infrastructure variability, and operational distances introduce additional complexity. Exploration of a distributed infrastructure architecture analogous to the Trans-European Transport Network, adapted for Indo-Pacific conditions, would provide valuable insight into large-scale network design.

Modeling efforts should incorporate dynamic adversary behavior and real-time disruption scenarios to better capture the evolution of infrastructure dependency under contested conditions. Integration of industrial base data (including production capacity, supply chain dependencies, and labor constraints) would improve the fidelity of both dependency assessment and procurement simulations.

Additional experimentation should apply the framework across a broader range of capability categories, including unmanned systems, logistics throughput enablers, and sustainment capabilities. Comparative analysis across these categories would help validate the generalizability of the framework and identify category-specific dynamics.

Field-based validation is also required. The framework should be incorporated into exercises, war games, and operational planning efforts to evaluate its effectiveness under real-world conditions, including the integration of dependency analysis into planning processes, the development of functional requirements, and the responsiveness of procurement mechanisms.

Future research should also advance program-based or web-enabled implementations of GEO-LENS and VPC that leverage artificial intelligence and Internet of Things (IoT)-enabled data streams to enable real-time or near-real-time assessment and procurement adaptation. These systems would ingest live data from infrastructure, logistics networks, energy systems, contractor production lines, and commercial supply chains to continuously update Operational Dependency Index scoring, dynamically refine requirements, and adjust contract allocation across vendors based on evolving cost, schedule, and performance conditions.

AI-enabled analytics and machine learning models could identify emerging constraint patterns, forecast disruption risk, and recommend allocation shifts within the VPC construct, transforming the framework from a periodic analytical tool into a continuously adaptive decision system. Future modeling and field experimentation should evaluate the feasibility, security, and



decision advantage of these architectures within contested environments, including integration with joint command-and-control systems, resilience under degraded or denied communications, and the implications of automated or human-in-the-loop contract reallocation.

Finally, pilot programs should be conducted to test portfolio-based procurement structures and AI-enabled GEO-LENS/VPC implementations within operational acquisition environments. These pilots would provide empirical data on contractor behavior, performance outcomes, supply resilience, and the effectiveness of real-time adaptive decision systems, enabling continuous alignment between operational dependency, requirements, and procurement execution at operationally relevant speeds.

## **Recommendations for Action**

The findings of this research have immediate implications for the military, the acquisition community, the defense industrial base, and the research community. While long-term reform will require sustained effort, several near-term actions can be taken to address the identified challenges.

Military organizations must integrate structured dependency analysis into operational planning and treat infrastructure as a pacing function rather than a supporting consideration. In contested environments, infrastructure performance governs the feasibility, tempo, and persistence of operations; failure to explicitly account for dependency risks results in plans that are optimized for conditions that will not exist in conflict. Exercises and war games must therefore incorporate infrastructure degradation, contested access, and sustainment constraints as primary variables. This requires accepting increased planning complexity and near-term reductions in perceived readiness in order to expose vulnerabilities early and avoid operational failure at scale.

Acquisition and procurement organizations must transition from static, single-award contracting toward portfolio-based procurement models capable of responding to distributed and evolving demand. Pilot programs should focus on high-quantity, scalable capabilities where competition can be sustained over time. Requirements development must prioritize functional performance, modularity, and scalability to enable continuous competition and adaptation. This shift requires accepting near-term inefficiencies, reduced predictability in contract outcomes, and increased management complexity in order to reduce long-term supply risk, accelerate innovation, and generate resilient, surge-capable industrial capacity.

The defense industrial base must transition from efficiency-optimized production models to flexible, surge-capable systems designed to respond to distributed and evolving demand signals generated by operational dependency. Under competitive procurement structures, firms no longer compete to win contracts; they compete to retain and expand production share. Allocation is continuously adjusted based on cost, schedule, and performance outcomes, creating a selection environment in which adaptability determines survival. Firms that fail to invest in scalable production, resilient supply chains, and rapid reconfiguration capability will lose relevance as demand shifts. In this system, adaptability is not an advantage—it is a requirement for participation.

The academic and research community must expand interdisciplinary efforts linking operational analysis, infrastructure systems, supply chains, and industrial economics to support the development of analytically rigorous and operationally relevant frameworks. This includes advancing modeling approaches, validation methods, and data integration techniques that can capture dynamic, contested environments. Such work requires moving beyond domain-specific optimization toward integrated system analysis, accepting increased methodological complexity and uncertainty in order to better represent real-world conditions and inform decision-making under constraint.



## References

- Attar, M., Ayers, J., & Mooney, N. (2025, July). *Revitalizing national security supply chains to overcome today's unprecedented logistics challenges*. The MITRE Corporation. <https://www.mitre.org/sites/default/files/2025-07/PR-25-1988-Revitalizing-National-Security-Supply-Chains.pdf>.
- Austin, L. (2025). *Investment strategy for the Office of Strategic Capital*. U.S. Department of Defense.
- Caudle, D. (2026, February). *U. S. Navy fighting instructions*.
- Hegseth, P. (2026, January). *Department of War 2026 national defense strategy (NDS)*.
- Hughes, Z. (2024). Giving our “paper tiger” real teeth: Fixing the US military’s plans for contested logistics against China. *Joint Force Quarterly*, 115(3), 28–45. <https://www.ndu.edu/News/Article-View/Article/3960834/giving-our-paper-tiger-real-teeth-fixing-the-us-militarys-plans-for-contested-l/>.
- LaPlante, W., & Lowman, C. (2024, May). *2024 regional sustainment framework (RSF)*. Under Secretary of Defense (Acquisition & Sustainment)
- Pierce, A. (2025). Employing a variable portfolio contract model to accelerate innovation incorporation, enhance operational sustainability, and reduce supply risk in the procurement process. *22nd Naval Postgraduate School Acquisition Research Symposium*. <https://dair.nps.edu/handle/123456789/5422>.
- Potter, J., Gering, B., and Kilby, J. (2025a, December). *Portfolio acquisition executive for mission systems transition*. Assistant Secretary of the Navy (Research, Development, and Acquisition).
- Potter, J., Gering, B., and Kilby, J. (2025b, December). *Portfolio acquisition executive for munitions transition*. Assistant Secretary of the Navy (Research, Development, and Acquisition).
- Publications Office of the European Union. (2024, June). *Regulation (EU) 2024/1679 of the European Parliament and of the Council on Union Guidelines for the development of the Trans-European Transport Network (TEN-T)*.
- Reitter, N. (2026, March). *Advanced naval basing innovation forum: Strategic synthesis of expeditionary capabilities*. Navy Expeditionary Warfighting Development Center.
- Smith, E. (2024, August). *39th commandant's planning guidance*.
- Trump, D. (2025, November). *National security strategy (NSS) of the United States of America*.



# Demo to Deployment in 30 Days: The ARMS Case Study

**Matthew Cole**—is a Technical Project Manager at NSWC PHD, with a portfolio focusing on rapid development and readiness technology. He has a BS in Aerospace Engineering from MIT (2015) and an MA in Defense and Strategic Studies from the U.S. Naval War College (2024). From January 2025 to present, he served as the Project Manager for the Augmented Reality Maintenance System (ARMS). [matthew.d.cole14.civ@us.navy.mil]

**Nicholas Bernstein**—is a Research and Development Engineer in NSWC PHD's Office of Engineering. His background in neuroscience, biomedical engineering, computer science, and systems engineering has informed contributions to various Navy projects involving data analysis, augmented reality, and distance support. From 2022 to present, he served as NSWC PHD's Augmented Reality Maintenance System (ARMS) Engineering Lead, working with NAWCAD Lakehurst to mature the system towards deployment readiness. [nicholas.r.bernstein.civ@us.navy.mil]

## Abstract

In February 2025, Augmented Reality Maintenance System (ARMS) had a shipboard demonstration of its remote technical assistance capabilities. Video of the demonstration made its way to critical leaders across Systems Commands (SYSCOMs), and the ARMS team at Naval Surface Warfare Center Port Hueneme Division (NSWC PHD) was asked to field the system across a Carrier Strike Group within 30 days. Over the next month, the ARMS team overcame every programmatic and policy barrier to deploy an integrated, accredited, and capable system into the hands of the warfighter. This achievement was made possible through a variety of prototype and technology acceleration authorities, including creative contracting strategies, Non-Permanent Change (NPC) alteration status, the Fleet Risk Acceptance process, and an Urgent Deployment Test (UDT) Interim Authority to Test (IATT). The volume of necessary workarounds reveals several insights about technology transition in the U.S. Navy; namely, criticality of high-level advocacy, relationships with key Fleet staff, shipboard integration expertise, emergent labor funding, and flexible procurement options. This paper offers a case study of the ARMS rapid deployment effort, intended as a resource for acquisition professionals and an example of how a program can successfully innovate by trading risk for speed.

## Introduction

In the modern security ecosystem, “innovation” is frequently presented as both a watchword and an unalloyed good. The Trump administration references innovation as a core national objective in both the National Security Strategy (NSS; The White House, 2025) and the National Defense Strategy (DoW, 2026), as well as several key executive orders<sup>1</sup>. And for all the changes that the Trump administration has brought to the Department formerly known as Defense, this topic is notable for its continuity with the Biden Administration, which spoke of innovation in no less glowing terms in the former 2022 NSS (The White House, 2022).

To the national security establishment, innovation is primarily a matter of speed. Relatively modest increments of capability can be considered “innovative” if they are delivered fast enough to make a difference to the mission, and transformative, generation-defining programs are frequently considered innovation failures when they are delivered slowly or late. Acquisition professionals accustomed to making trades in the cost-schedule-performance decision space will recognize this paradigm as one in which schedule reigns supreme. But understanding the need to prioritize schedule in theory and achieving acquisition speed in practice can be quite different, owing to the wide breadth and depth of pitfalls that await the

---

<sup>1</sup> See for example, “Removing Barriers to American Leadership in Artificial Intelligence” (Exec. Order No. 14179) and “Modernizing Defense Acquisitions and Spurring Innovation in the Defense Industrial Base” (Exec. Order No. 14265).





unwary Program Manager, especially in the last tactical mile of delivering capability to the Fleet. To paraphrase Clausewitz, everything is very simple in deployment, but even the simplest thing is hard.

To assist Program Managers in that endeavor, this paper presents the case study of a single, highly successful attempt to deliver a modest capability to the fleet at maximum speed: the Augmented Reality Maintenance System (ARMS), which successfully accomplished an emergent, leadership-directed deployment across a battlegroup approximately 30 days after a successful technical demonstration.

There are limits to the value of any single case study in acquisition, because every acquisition system has a unique technical and programmatic posture. ARMS is no different, and many of the barriers and solutions applicable to ARMS may not extend to other systems. But by offering a complete, end-to-end account of how each of these barriers was identified and resolved, it is our hope that other acquisition professionals will be better equipped to address similar barriers in the future, even if they are different in the specifics.

This case study will be structured in three parts. In the first part, a basic timeline and description of activity will be provided, illustrating the history of ARMS and the key events over the course of the 30-day ARMS transition window. The second part will analyze the timeline by identifying what we view as the primary barriers, the solutions that were employed to resolve those barriers, and the key enabling factors that facilitated those solutions. Finally, the third part will provide discussion of the case study, crystalizing what we offer as the most salient lessons and recommendations for practitioners.

## **Case Narrative**

### **Inception To Technology Demonstration**

The fundamental problem that ARMS was developed to solve is the extreme delay and expense associated with sending Subject Matter Experts (SMEs) in person to deployed U.S. Navy ships, a process known as “On-Board Technical Assistance” or OBTA. The goal of the development work that culminated in ARMS was to offer an alternative that improved upon traditional distance support by phone and email, something that we have previously described as Remote Technical Assistance (RTA; Cole et al., 2025).

Initial development of the technology that would become ARMS was slow, steady, and sometimes meandering. Development, testing, and maturation of the tool on a technical level was a multi-year effort across Navy System Commands (SYSCOMs) with multiple failures and significant changes. Development of the core software of ARMS began in 2018 by Naval Air Warfare Center Aircraft Division (NAWC AD) Lakehurst’s (LKE) Collaborative Research, Engineering, Analysis, and Training in Immersive Virtual Environments (CREATIVE) Lab as part of a Naval Innovative Science and Engineering (NISE) challenge. Then named the Collaborative Augmented Reality Maintenance Aid (CARMA), the tool streamed audio and video between a Windows laptop and a standalone augmented reality (AR) headset. The chosen AR headset was a self-contained device that provided all the power, computation, and user interface necessary for a deployed maintainer to broadcast and receive real-time data streams with a remote SME. The SME, operating the software on a standard Navy asset, could advise the sailor while being virtually co-located in the same room.

In 2019, on the other side of the country and in another SYSCOM, NSWC PHD sought technology that could benefit the sustainment community with current logistical difficulties impacting distance support today and into the future. Many In-Service Engineering Agents (ISEAs) spend much of their time traveling to perform OBTA visits around the world. The following year, the COVID-19 pandemic greatly restricted travel of SMEs, slowing necessary



OBTAs globally. LKE's distance support tool, now called Augmented Reality Remote Maintenance Support Service (ARRMSS), participated in the 2021 Advanced Naval Technology Exercise (ANTX)–Coastal Trident exercise; the distance support capability used PHD's self-defense test ship as a demonstrative underway platform for a proof of concept ship-to-shore remote expert call. What would soon be renamed once more to ARMS proved the capability for remote users to communicate between lab sites and across miles of ocean.

Many years earlier, NSWC PHD had prototyped a similar real-time distance support tool which, while mobile, had a much larger footprint; the bulky Distance Communications Maintenance System (DCoMS) provided a similarly valuable capability as ARMS but required the sailor to carry a backpack holding approximately fifty pounds of gear. Recognizing the need but the significant advancement ARMS represented due to reduction in size for the distance support concept, NSWC PHD quickly partnered with LKE to mature and field the system. Soon after, Program Executive Office Integrated Warfare Systems (PEO IWS) Program Executive Officer RDML Seiko Okano, aware of the potential impact of remote technical assistance in remote regions, provided funding to deploy the tool on active Navy platforms.

The next few years led to several shifts in technological approach as the realities of system integration came into full view. Despite the fact ARMS was built with Navy use cases and network considerations in mind, the chosen AR headset proved incapable of fully integrating with Navy shipboard networks or fully complying with all relevant Security Technical Implementation Guides (STIGs), a requirement for authority to operate (ATO) approval (DoD, 2022). Several lab-based tests and fleet exercises were successful to certain degrees, but full shipboard integration on active Navy ships and networks was invariably blocked from full functionality for one reason or another. The AR headset's operating system was simply unable to pass cybersecurity muster. Further, no commercial headset alternatives had the features required or developmental longevity to warrant long-term technological investment. As a result, the team continued to test the headset ARMS version while concurrently developing a laptop-based alternative using smart glasses peripherals instead of a self-contained AR device.

Despite these barriers to fielding the form factor, the AR headset version of ARMS was installed on three ships for short-duration testing using an interim authority to test (IATT) and limited ship change document (SCD) for a non-permanent change (NPC). One ARMS demonstration performed on the USS Spruance was recorded and passed on to then-RADM Seiko Okano, now Commander of Naval Information Warfare Systems Command (NAVWAR). RADM Okano was concerned about NAVWAR's ability to support the USS Nimitz and accompanying escort Destroyers (collectively "Carrier Strike Group" or "CSG") during that ship's final deployment prior to expected decommissioning. Seeing ARMS as a potential lifeline for the CSG, RADM Okano formally requested that NSWC PHD deploy an operational variant of ARMS across all the ships in the battlegroup prior to their deployment, roughly 30 days later.

### **Technology Demonstration to System Deployment**

There was significant debate within NSWC PHD about whether the timeline was feasible. While the underlying software had a lengthy development history, the operational, ship-ready system that RADM Okano requested simply did not exist. The hypothesized laptop-based replacement configuration had never been assembled in sufficient quantity to outfit an entire CSG. Even if it had, that new configuration would lack the necessary cybersecurity approval to authorize its use on a ship, and the shore-based infrastructure needed to consummate the ship connection was due to be shut down within the month. And even ignoring all technical concerns, the procedural requirements associated with authorizing and accomplishing an installation on a U.S. Navy ship alone would have made the timeline tight enough to raise eyebrows. The decision to proceed with the attempt, directly authorized by NSWC PHD Commanding Officer



CAPT Anthony Holmes, was made with the acceptance of risk and the knowledge that success was by no means a forgone conclusion.

The first and most immediate challenge was the fact that there was no budget to support this emergent, unfunded requirement. This challenge was resolved through the immediate reprioritization of NSWC PHD NISE funds. Because of its flexibility, NISE could be used to fund government labor across the warfare center (including, critically, travel and overtime), allowing for the immediate activation of a half dozen branches with critical functions including cybersecurity, installation oversight, availability planning, configuration management, and technical documentation. NSWC PHD's primary role as the ISEA for most surface Navy combat systems was crucial in this regard. In business terms, NSWC PHD possessed an organic, vertically integrated capacity to perform almost all parts of the technology deployment process for in-service ships. The combination of a flexible funding source with this in-house vertical integration was one of the decisive factors in enabling rapid ARMS fielding.

The second immediate challenge was coordination across this wide array of disparate branches. This was accomplished through two means: first, the centralization of project management authority under a single Project Manager (PM), and second, a remarkably flat organizational structure, depicted in Figure 1. Daily, high-level project meetings provided a consistent drumbeat of updates, allowing for maximally dynamic schedule updates.




**Figure 1. ARMS Team Organizational Chart**

This highly central, highly flat organizational structure allowed for parallel execution of the many subtasks needed to execute ARMS fielding across the USS Nimitz CSG. The major subtasks and the relationships between them are depicted in Figure 2.




**Figure 2. ARMS Fielding Tasks and Relationships**

The single most critical and longest lead subtasks respectively were the development of the system baseline and the procurement of required parts. These two tasks were related, because the system baseline consists of both the hardware and the software components, and to the extent that the team needed to make hardware substitutions to account for delivery timelines, those substitutions would have to be determined before the baseline could be finalized. Procurement was an especially steep challenge, because the deployment timeline precluded the use of even the fastest Government mechanisms, like Purchase Cards and Defense Logistics Agency (DLA) Troop Support. Fortunately, NSW PHD had a flexible contract in place with a vendor that had been pre-funded with NISE dollars in anticipation of emergent needs, and this allowed for the acquisition of short-lead material on a much-accelerated timeline. But even the vendor could not accelerate the delivery of the long-lead laptops. To acquire these assets, NSW PHD requested assistance from Program Management Aviation (PMA) 260, a Naval Air Systems Command (NAVAIR) program office that fields a similar device as a preventative maintenance aid for U.S. Navy aircraft maintainers. PMA 260 generously provided enough units from their existing stock to cover the USS Nimitz CSG deployment.

The software baseline was also a challenge to finalize. The ARMS client software was finalized and tested, but to deploy a full-up-round system, it would need to be loaded onto a hardened Operating System (OS) capable of passing the rigorous battery of cybersecurity tests needed for an operational IATT. After considering several options, the team elected to build the OS from the ground up, leveraging a Secure Host Baseline (SHB) previously maintained by the Defense Information Security Administration (DISA). This resulted in a stand-alone system stripped down to the bare minimum software needed to function, reducing cybersecurity risk at the cost of functionality. While not necessarily ideal, this compromise was deemed to present less schedule risk than the alternative approach, which could have extended the IATT approval timeline with additional rounds of review and correction.

To achieve the required IATT approval within the 30-day deployment window, the team relied on a new type of IATT specifically intended to support rapid fielding, called an Urgent Deployment Test (UDT). This type of IATT requires that the system be used in support of an

operational unit rather than purely for research and development and requires a written endorsement from a flag-level officer. RADM Okano provided the required endorsement for ARMS, making the deployment timeline at least feasible. But it still took a heroic effort from the NSWC PHD cybersecurity team and in-house validators to create and submit the IATT package in the required system to route it for approval. Even then, full approval was not obtained until the very final hour of the ARMS installation window, almost literally.

Once the ARMS system baseline was established, the next most critical task was network integration. For ARMS to be usable on the shipboard network present on U.S. Navy carriers and destroyers, it needed to be certified for use with Consolidated Afloat Networks and Enterprise Services (CANES). This testing was previously done for the ARMS headset configuration, but there was no guarantee that the old routing rules would still work with the new configuration without repeat testing. Fortunately, in the run-up to the rapid deployment window, the team was already working through the integration lab's approval chain to test the laptop-based alternative ARMS kit. Once in the lab, the NAVWAR team quickly applied previously established network configuration changes developed in the lead-up to the headset demonstration and confirmed operation of the new laptop-based ARMS kit, eliminating the single most important technical barrier to deployment.

To obtain authorization to install ARMS on US Navy ships outside of a designated availability window, the team chose to pursue the Fleet Risk Acceptance (RA) process. This process allows Fleet Commanders to authorize ship alterations that do not meet the traditional milestones for maturity dictated by the full Navy Modernization Process (NMP; Naval Sea Systems Command, 2021). The RA process is quite powerful, but it does present a few hard requirements. For one, the alteration seeking authorization must at least be documented by a SCD, even without maturity. It must also have an Authority to Operate (ATO) or IATT, there must be sufficient evidence that it will not negatively affect the network or the ship, and even the RA process does not exempt an alteration from complying with production requirements enforced by the Regional Maintenance and Modernization Coordination Office (RMMCO). But perhaps the most important prerequisites to effective use of the RA process are relationships and communication. If Type Commander and Fleet staff are not aware of the alteration, the justification, and the urgency, requests for RA are unlikely to succeed.

Accomplishing this constellation of requirements took the coordinated work of several teams. An NSWC PHD configuration management team built an SCD with inputs from engineering, while a technical documentation team worked with engineering to build a technical manual, training materials, and installation and verification procedures needed for the Quality Assurance (QA) workbook, a RMMCO requirement. With the SCD and supporting documentation in hand, an Availability Planning team handled the Naval Message traffic needed to initiate the RA process, and NSWC PHD embedded representatives at the Type Commander and Fleet staff were instrumental in socializing the urgency of the alteration. Collectively, we estimate that approximately two dozen Navy civilians and a handful of uniformed officers were directly involved in creating, reviewing, routing, and approving the paperwork required to authorize the ARMS installation.

The final prerequisite for accomplishing installation within the 30-day window was the necessity of setting up new shore-side infrastructure. The requirement from RADM Okano and CAPT Holmes was that the infrastructure support access by any relevant shore SME, including personnel at NAVWAR's facilities on both coasts, multiple Warfare Centers, and Regional Maintenance Centers (RMCs), including those in forward deployed locations, who have a critical role as first responders to Fleet technical issues. While the specific technical details employed by the ARMS team to meet these requirements are beyond the scope of this paper, the approach broadly involved leveraging the Research, Development, Test, and Evaluation (RDTE)





networks present at and owned by NSWC PHD. By hosting certain components locally and utilizing peer network connections to remote sites with known and trusted protocols, the team devised a method for arbitrary SMEs to access ARMS without a requirement for locally installed software.

Finally, in the last days of March 2025, these many lines of effort collided in the first ever installation of the operational ARMS configuration aboard an active U.S. Navy ship. As a testament to the technical work, planning, and preparation completed in this short period, the team was able to successfully install the system of two Destroyers in a single day. By the end of the month, every ship in the CSG had been equipped with ARMS and tested satisfactory (apart from a leading deployer that had departed weeks sooner).

## Post Deployment Developments

The reward for success in installing across the Nimitz CSG was a request to repeat the effort multiple times. In the year that followed, ARMS kits were built and installed on multiple carrier strike groups, amphibious ready groups, forward deployed ships, and ashore platforms. Approximately half the surface fleet—including destroyers, aircraft carriers, and amphibious ships—are now equipped with ARMS with more installations planned. Hundreds of shoreside SME accounts have been created. At time of writing, ARMS is estimated to have prevented 85.5 combined days of system downtime and saved 3,600 SME travel hours. The reduced SME travel and downtime have saved the Navy over \$600,000. Even when ARMS cannot solve an issue through exclusively remote SME guidance, the advanced information allows the SME to travel with more information and the correct equipment needed to correct the issue. What's more, ARMS is challenging how maintenance is prioritized and performed across the organization because the system SME no longer needs to be physically present at the nearest port.

## Analysis

Based on an analysis of the provided timeline, the experiences of the authors leading and executing the ARMS transition, and interviews with other personnel involved in the process, we have identified seven primary barriers that the ARMS program encountered and overcame during the 30-day Nimitz CSG deployment. These seven barriers are identified in Table 1.

**Table 1. ARMS Barriers and Solutions**

| Barrier                              | Solution                                                      | Enablers                                                              |
|--------------------------------------|---------------------------------------------------------------|-----------------------------------------------------------------------|
| 1. Urgent Labor Funding Requirements | Naval Innovative Science and Engineering (NISE) Funds         | Leadership prioritization                                             |
| 2. Rapid Technology Adaptation       | Standalone configuration                                      | Available baseline, existing hardware BOM                             |
| 3. Shipboard Network Integration     | Network lab time, available network SMEs                      | DICE Lab, High level advocacy                                         |
| 4. Authority to Operate (ATO)        | Urgent Deployment Test (UDT) Interim Authority To Test (IATT) | Flag endorsement, ISSE and validator support                          |
| 5. Procurement                       | Flexible prepositioned funds, cross-program asset borrowing   | NISE, PMA 260 support                                                 |
| 6. Alteration Authorization          | Non-Permanent Change (NPC) process, Emergent Change Process   | Organic warfare center expertise and relationships with fleet and PMO |
| 7. Production Requirements           | Government AITM and OSIC support                              | Organic warfare center expertise                                      |



## **Urgent Labor Funding Requirements**

To rapidly task the many personnel across several technical areas, the ARMS project needed access to a flexible source of funding for labor, travel, and overtime. This came in the form of NISE funds, a specialized type of overhead funding available to Department of Defense laboratories, including Naval Warfare Centers (10 U.S.C. § 4123). While NISE-funded projects typically run on an 18-month cycle from proposal to approval to execution, within NAVSEA, local control of NISE funding is sufficient to enable some degree of mid-year reprioritization based on emergent needs. This flexibility was exactly what the ARMS project needed at the crucial moment it was tasked to deliver. The key enabling resource to unlock NISE funds for this purpose was leadership prioritization, in this case that of CAPT Anthony Holmes, the NSWC PHD Commanding Officer.

## **Rapid Technology Adaptation**

To adapt the effective but insecure pilot configuration into a secure, operational ARMS system, the team had to re-evaluate what form factor could best deliver the ARMS software while meeting security requirements. The critical breakthrough was the decision to build a stand-alone laptop configuration with a Defense Information Systems Agency (DISA) secure host baseline (SHB) for the Windows installation. Further, this kit included a small carrying case, headphones, webcams, and smart glasses. The smart glasses provided some limited AR capability but primarily provided a hands-free display for the maintainer and a camera feed back to the remote expert. Concurrently with the ongoing shipboard testing of the AR headset leading up to the USS Spruance demo, this known-good hardware was conceived and assembled to support shipboard network integration testing.

The key enabling resources for this solution were the fact that the team had already investigated and developed a notional Bill of Materials (BOM) for a laptop-based kit prior to tasking and the availability of a suitable SHB. This latter enabler is notable because with the sunset of the Windows 10 operating system, available SHBs will no longer meet cybersecurity requirements. Future systems may have more difficulty rapidly meeting cybersecurity requirements without inheriting controls from the host network, increasing integration complexity.

## **Shipboard Network Integration**

Given the reliance of modern warships and systems on advanced IP-based internal networking, network integration is often the single most important technical barrier for any shipboard system that requires a direct IP-based connection, and network certification is one of the most important non-technical barriers. In general, the lead time to enter the simulated shipboard network lab and certify a new system can be nearly a year. The ARMS team was able to complete lab-based certification testing within 30-day deployment window only because of a few decisive factors: an in-progress lab entry request, previous relationships, and established technical artifacts.

First, at the start of the 30-day window, the request to test alternative laptop-based kit was already submitted and reviewed at several layers of the intake process; the system was, despite a change in hardware, largely the same in how it interfaced with the shipboard network and represented a fairly small hardware deviation for lab cyber reviewers. Secondly, the team knew who could assist in expediting the final review stages of the lab entry request and expressed the urgency of the installation schedule; although no cyber scrutiny was bypassed, the requests weren't subject to the usual review windows. Finally, because the team had spent several weeks spread over the previous few years in the integration lab testing multiple ARMS iterations, the bulk of configuration changes necessary for the laptop-based kit were known. By learning what didn't work from past attempts, the team was able to make the most of allocated lab time. Additionally, the history of ARMS troubleshooting leading up to the successful USS



Spruance demo using the AR headset informed the downstream network configuration changes necessary for proper data transport.

The key enabling resources that facilitated this solution were the availability of the network test lab during the period required and the availability of high-level advocacy to prioritize ARMS for that lab time.

### **Authority to Operate**

The process to obtain an Authority to Operate (ATO) within the Department of the Navy's Risk Management Framework (RMF), including an IATT, is a notoriously lengthy process. In fact, this barrier would have single-handedly rendered the Nimitz deployment timeline infeasible if not for the availability of a specialized, expedited process called an Urgent Deployment Test (UDT). As described above, this IATT variation requires that the system be intended for operational use, and that it receive endorsement from a flag-level officer or civilian equivalent. For systems capable of meeting those bars, the UDT IATT may be an appealing choice to facilitate rapid deployments. To expedite this variation, the ARMS team socialized the request early with Information Systems Security Engineers (ISSE) and key validators.

### **Procurement**

When Department of the Navy (DoN) programs require procurement on accelerated timelines, there are several feasible options that offer significantly greater speed than traditional contracts under the Federal Acquisition Regulations, including Government Purchase Cards and the Defense Logistics Agency (DLA) Troop Support. But the ARMS project found that all these options typically result in material delivery approximately 30–60 days after the identification of need after accounting for administrative delay. Options for procurement when material is needed in the 0–30 day window are significantly more limited. Fortunately, ARMS was able to leverage an existing contract that NSWC PHD held with a vendor for rapid prototyping support. Using prepositioned NISE funds, this contract allowed the vendor to procure the needed components on NSWC PHD's behalf. This contract proved to be fast enough to procure the needed short-lead kit components in time for deployment.

For the long-lead laptops, other methods were required. In the end, the team determined that the only feasible method to acquire enough suitable, ruggedized laptops was to borrow them from another DoN partner. These were graciously provided from the stock of PMA 260, NAVAIR's program office for Common Aviation Support Equipment. Support from PMA 260 was the critical enabler that allowed the ARMS team to obtain the material needed to outfit the Nimitz CSG.

### **Alteration Authorization**

As a matter of normal business under the Navy Modernization Program (NMP), alterations to in-service ships are documented by SCDs and performed during Selected Restricted Availabilities (SRAs). This process is documented in the NMP Management and Operations Manual (NMP-MOM), which spans 943 pages and includes 9 appendices. For urgent, safety or mission-driven alterations requiring installation in 45 days or less, the Emergent Change Process offers a much faster alternative that allows the alteration to be matured in parallel with (in practice, typically after) execution of the alteration. The emergent change process was the primary procedural vehicle that ARMS leveraged to authorize installation within the 30-day Nimitz CSG install window, in conjunction with a Non-Permanent Change (NPC) SCD.

To overcome this barrier, the ARMS team brought to bear two key enablers: organic expertise at NSWC PHD in the NMP and existing relationships with Type Commander and Fleet Commander staff. The fleet relationships were particularly important, because endorsement and



authorization from these actors are required for the emergent change process. Early socialization with the fleet is key to achieving risk acceptance and authorization.

### **Production Requirements**

To execute on an in-service ship, all alterations regardless of authorization mechanism must be approved by the Regional Maintenance and Modernization Coordination Office (RMMCO). RMMCO acts as a gatekeeper to ensure that the teams performing alterations meet the requirements to execute safely and in accordance with U.S. Navy standards, with appropriate oversight. Most RMMCO requirements are met by compiling a QA Workbook, typically a physical binder that contains all the documents, personnel certifications, signature pages, references, and other information needed by RMMCO.

The ARMS team was able to overcome this barrier by leveraging organic capability within NSWC PHD, specifically Alteration/Installation Team Manager (AITM) and On-Site Installation Coordinator (OSIC) support. Because NSWC PHD performs a large volume of ship alterations (over 400 annually), the team was able to rapidly build the required QA workbook from existing templates, and leveraged existing, certified, and experienced waterfront personnel.

### **Discussion and Recommendations**

#### **The Importance of Maintaining Flexibility in Funding**

Because government agencies are under significant pressure to expend all their funds within strict periods of performance, there is a tendency by many organizations to plan often and early, even for funding types that would theoretically offer greater flexibility. While this surplus of planning reduces the risk of under- or over-spending, it also comes at the cost of limiting the ability for mid-period adaptation to support innovation and rapid technology transition. In the case of ARMS, incurring the risk of a mid-year NISE replan was a critical and necessary cost of responding rapidly to an emergent need for technology transition. One takeaway for acquisition professionals is that in limited and responsible ways, it may be prudent to actively avoid the temptation to over-plan or at least identify the slack available within the plan.

#### **The Limitations of “Rapid” Procurement**

The experience of ARMS points to a critical gap in options for programs requiring procurement of material in less than 30 days. By process optimization and an intense focus on reducing administrative delays to the minimum possible, it might be possible to shrink this window to perhaps 21 days. But the key takeaway for acquisition professionals is that any procurement needed on a “hyper-rapid” timeline simply cannot afford to start from scratch. The obvious solution is a specialized procurement contract with a vendor and prepositioned funding, similar to what ARMS leveraged for short lead items. But the best and sometimes only fallback strategy for a program manager in a pinch may be to phone another program with material already on hand. For individual acquisition professionals, this underscores the value of a robust professional network. For the enterprise, this underscores the value of shared data visibility to better equip the acquisition community with information about who may be able to help meet urgent critical needs.

#### **The Criticality of Relationships and Advocacy**

While so-called “bottom up” innovation that starts as a good idea at the working level can be a powerful source of transformation within an enterprise over time, it is worth noting that the ARMS team’s rapid success was driven largely by the ability to leverage top-down support. RADM Okano’s request to deploy for the Nimitz CSG functioned as a critical stimulus to kick start the rapid transition, and Capt Holmes’s support unlocked the full potential of the warfare center tasked with executing it. RADM Okano also provided the required flag endorsement



without which it would have been impossible for ARMS to receive an IATT, and as the Functional Authorizing Official (FAO), also served as signature authority on the IATT itself.

Fleet relationships were also essential for navigating the emergent change process, and for the coordination of execution logistics, there is no effective substitute for direct waterfront relationships with ships' crews. These relationships are difficult to forge in a hurry, and the success of the ARMS team owed primarily to the existing organic capabilities of NSWC PHD. For Acquisition professionals looking to break into this space, a transition partner with established fleet relationships may be a powerful asset.

## **Risks and Waivers**

It is no exaggeration to suggest that rapid transition is predominantly an exercise in the acceptance of risk. The primary tool for accepting risk within the U.S. Navy is typically a "waiver" from the usual, lengthy process, which may take various forms depending on the process in question. ARMS used many waivers to accomplish rapid deployment, waiving requirements that varied from logistics product requirements to cyber controls to installation dates.

Waivers are a powerful and sometimes necessary tool, but they do have an important downside: the typical process always exists for a reason. Whether for safety, or cybersecurity, or interoperability, all the laborious layers of typical Navy review do important jobs for the fleet. Waiving a process, while expedient, also deprives the DoN of all the benefits of those layers of review. We suggest that in some areas, there may be additional opportunities to find a middle ground, allowing for faster outcomes without relinquishing all the benefits of full review.

## **Warfare Centers as Rapid Transition Agents**

Naval Warfare Centers are Navy Working Capital Fund (NWCF) organizations that provide organic government capability across many functions for a variety of customers. They are considered Department of Defense Laboratories (10 U.S.C. § 4121) and have a mission that spans from Research and Development (R&D) to in-service fleet support. Warfare Centers function like internal contractors within the DoN, and their core work is to accomplish tasks that are either inherently governmental functions or not cost effective for external contractors to provide as a service (Naval Sea Systems Command, n.d.). One important distinguishing feature of Warfare Centers is that they only act as service providers, and do not directly receive congressional appropriations or have their own requirements.

As the example of ARMS illustrates, Warfare Centers bring a powerful set of tools under one roof, including integration and mission engineering, logistics, configuration management, technical documentation production, cybersecurity engineering, alteration/installation management and oversight, and project management. Consequently, Warfare Centers are uniquely well-positioned to help accelerate technology transition, with one important caveat: they cannot do it on their own initiative. It takes an external requirement (and frequently external funding) to stimulate a Warfare Center to assume the role of a rapid transition agent. Even then, there are limitations, because Warfare Centers may not be well-positioned to do some of the foundational engineering needed to mature technologies at lower readiness levels. But for acquisition professionals tasked with fielding already mature technologies to the fleet at maximum speed, Warfare Centers may prove to be especially valuable partners.

## **Conclusion**

The challenges of program management within the DoN often give rise to a peculiar variant of the so-called Dunning-Kruger effect. Those who have never attempted to field a system to the fleet generally underestimate the difficulty and effort of doing so at speed. But by the same token, those who have tried and failed often succumb to a form of bureaucratic





nihilism and assume that truly rapid deployment is impossible. The case study of ARMS refutes both biases.

To the uninitiated, the challenge of taking a system that accomplished a successful technology demonstration on one ship and deploying it to a handful more may seem trivial. The many technical and non-technical barriers ARMS confronted and overcame shows that this is not the case. Even the simplest, most mature system can expect to encounter similar roadblocks, often requiring the expertise of a wide array of experienced professionals to effectively resolve. Any acquisition professional attempting rapid technology transition should be clear eyed about the difficulty and cost associated with delivering at speed.

But the notion that processes like NMP or cybersecurity authorization simply cannot be accomplished quickly is outdated and wrong. As ARMS shows, with the right approaches and enabling factors, these barriers are not unmanageable. In particular, funding flexibility, high level advocacy, mature relationships with fleet stakeholders and other acquisition professionals, and responsible use of process waivers are key factors that can reduce friction and speed deployment.

Every program is always unique in the specifics. The barriers and solutions employed by ARMS to facilitate a 30-day deployment may be irrelevant or infeasible for the next system to attempt deployment on a similar timeframe. The process of identifying and overcoming barriers, however, is one that all program managers must master. ARMS offers one successful model for acquisition professionals to consider.

## Appendix A: Acronyms and Abbreviations

| Acronym / Abbreviation | Definition                                                                                    |
|------------------------|-----------------------------------------------------------------------------------------------|
| AITM                   | Alteration/Installation Team Manager                                                          |
| ANTX                   | Advanced Naval Technology Exercise                                                            |
| AR                     | Augmented Reality                                                                             |
| ARMS                   | Augmented Reality Maintenance System                                                          |
| ARRMSS                 | Augmented Reality Remote Maintenance Support Service                                          |
| ATO                    | Authority To Operate                                                                          |
| BOM                    | Bill Of Material                                                                              |
| CANES                  | Consolidated Afloat Network and Enterprise Services                                           |
| CARMA                  | Collaborative Augmented Reality Maintenance Aid                                               |
| CREATIVE               | Collaborative Research, Engineering, Analysis, and Training in Immersive Virtual Environments |
| CSG                    | Carrier Strike Group                                                                          |
| DCoMS                  | Distance Communications Maintenance System                                                    |
| DISA                   | Defense Information Security Administration                                                   |
| DLA                    | Defense Logistics Agency                                                                      |
| DoN                    | Department of the Navy                                                                        |
| FAO                    | Functional Authorizing Official                                                               |
| IATT                   | Interim Authority To Test                                                                     |
| ISEA                   | In-Service Engineering Agent                                                                  |
| ISSE                   | Information Systems Security Engineer                                                         |
| LKE                    | Lakehurst                                                                                     |



|          |                                                               |
|----------|---------------------------------------------------------------|
| NAVAIR   | Naval Air Systems Command                                     |
| NAVWAR   | Naval Information Warfare Systems Command                     |
| NAWC AD  | Naval Air Warfare Center, Aircraft Division                   |
| NISE     | Naval Innovation Science and Engineering                      |
| NMP      | Navy Modernization Process                                    |
| NMP-MOM  | Navy Modernization Process - Management and Operations Manual |
| NPC      | Non-Permanent Change                                          |
| NSS      | National Security Strategy                                    |
| NSWC PHD | Naval Surface Warfare Center, Port Hueneme Division           |
| OBTA     | On-Board Technical Assistance                                 |
| OS       | Operating System                                              |
| OSIC     | On-Site Installation Coordinator                              |
| PM       | Project Manager                                               |
| PMA      | Program Management Aviation                                   |
| QA       | Quality Assurance                                             |
| R&D      | Research & Development                                        |
| RA       | Risk Acceptance                                               |
| RMC      | Regional Maintenance Center                                   |
| RMF      | Risk Management Framework                                     |
| RMMCO    | Regional Maintenance and Modernization Coordination Office    |
| RTA      | Remote Technical Assistance                                   |
| SCD      | Ship Change Document                                          |
| SHB      | Secure Host Baseline                                          |
| SME      | Subject Matter Expert                                         |
| STIG     | Security Technical Implementation Guides                      |
| SYSKOM   | System Command                                                |
| UDT      | Urgent Deployment Test                                        |

## References

- Cole, M., Bernstein, N., & Tenhoff, C. (2025). *Trends and advances in remote technical assistance (RTA) to improve fleet operational availability* [Paper presentation]. American Society of Naval Engineers 2025 Fleet Maintenance and Modernization Symposium. Available from the author upon request.
- Exec. Order No. 14179, 90 Fed. Reg. 8741 (2025). <https://www.federalregister.gov/d/2025-02172>
- Exec. Order No. 14265, 90 Fed. Reg. 15621 (2025). <https://www.federalregister.gov/d/2025-06461>
- Mechanisms to Provide Funds for Defense Laboratories and Test Organizations for Research and Development of Technologies for Military Missions, 10 U.S.C. § 4123 (2018).
- Naval Sea Systems Command. (2021, September 2). *Navy modernization process management and operations manual*, Revision 4, Change 2.



- Naval Sea Systems Command. (n.d.). *What we do*. Retrieved March 31, 2026, from <https://www.navsea.navy.mil/Home/Warfare-Centers/What-We-Do/>
- Science and Technology Reinvention Laboratories: Authority and Designation, 10 U.S.C. § 4121 (2018).
- DoD. (2022, July 19). *Risk Management Framework for DoD Systems* (DoD Instruction 8510.01). <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/851001p.pdf>
- DoW. (2026). *2026 national defense strategy*. <https://media.defense.gov/2026/Jan/23/2003864773/-1/-1/0/2026-NATIONAL-DEFENSE-STRATEGY.PDF>
- The White House. (2022). *2022 national security strategy*. <https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/10/Biden-Harris-Administrations-National-Security-Strategy-10.2022.pdf>
- The White House. (2025). *2025 national Security Strategy*. <https://www.whitehouse.gov/wp-content/uploads/2025/12/2025-National-Security-Strategy.pdf>











Naval  
Postgraduate  
School  
**Foundation**



ACQUISITION RESEARCH PROGRAM  
DEPARTMENT OF ACQUISITION, FINANCE, AND MANPOWER  
NAVAL POSTGRADUATE SCHOOL  
555 DYER ROAD, INGERSOLL HALL  
MONTEREY, CA 93943

[WWW.ACQUISITIONRESEARCH.NET](http://WWW.ACQUISITIONRESEARCH.NET)