SYM-AM-26-067

Proceedings
of the
Twenty-Third Annual Acquisition Research Symposium and Innovation Summit

Thursday, May 7, 2026 Sessions
Volume II
"Accelerating Warfighting Capabilities"

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

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
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/**

Table of Contents

WELCOME: DR. MICHAEL HESSE, VICE PROVOST FOR RESEARCH AND INNOVATION, NAVAL POSTGRADUATE SCHOOL

Dr. Michael Hesse—serves as Vice Provost for Research and Innovation (VPRI) at the Naval Postgraduate School (NPS), where he provides strategic leadership for the university's research enterprise. In this role, he oversees the development and execution of research priorities that advance the Navy, Marine Corps, and Department of Defense. He directs efforts to expand interdisciplinary collaboration, foster partnerships with government, industry, and academia, and support faculty and student innovation. Dr. Hesse also plays a central role in advancing the Naval Innovation Center at NPS, aligning cutting-edge research with operational challenges to accelerate solutions for national security.

Prior to coming to NPS, Dr. Hesse was the Director of the Science Directorate at NASA's Ames Research Center, a position in the Senior Executive Service. In this role, he was leading a staff of about 500 government and contract employees, and postdoctoral fellows engaged in Earth science, planetary and astrophysical research, and space biological research. His responsibilities included provision of an entrepreneurial vision, scientific leadership, the management of organizational resources, management and development of budget and financial resources, internal and external partnerships, and overall implementation of NASA priorities. He represented the organization to senior leadership of government and nongovernmental organizations inside the U.S. and internationally.

Before joining Ames, Dr. Hesse spent three years at the University of Bergen in Norway, where he held a professorship in physics. In Bergen, he also led the Geomagnetic Expert Service Centre, a multi-national consortium providing space weather services to ESA. He further held a part-time position with the Southwest Research Institute in San Antonio, where he contributed to research and leadership of NASA's Magnetospheric Multiscale mission.

Prior to moving to Bergen, Dr. Hesse had a distinguished 25-year career at NASA's Goddard Space Flight Center (GSFC), culminating in his role as the Director of the Heliophysics Science Division, a Senior Executive Service position. Dr. Hesse was also the founding Director of the Community Coordinated Modeling Center (CCMC), the world's foremost organization providing unprecedented comprehensive modeling services to the international space science community.

Dr. Hesse had various roles in NASA missions, most recently that of Lead Co-Investigator for Theory and Modeling for NASA's Magnetospheric Multiscale mission. He remains a publishing research scientist in space science and space weather, with more than 300 papers in the scientific literature, and an H index of 75. He was elected Fellow of the American Geophysical Union in 2010, he received NASA's Outstanding Leadership Medal in 2007, and NASA's Distinguished Service Medal in 2017. Dr. Hesse was elected member of Academia Europea in 2019, and he received the Space Weather and Nonlinear Waves and Processes prize of the American Geophysical Union in the same year. In 2023 he was recognized by the Meritorious Presidential Rank Award for Senior Executives. Dr. Hesse received his diploma and doctoral degree in Theoretical Physics from the Ruhr-Universität in Bochum, Germany. Dr. Hesse's hobbies are hiking, travel, sports cars, and amateur photography.

THIS PAGE INTENTIONALLY LEFT BLANK

WELCOME: HONORABLE MICHAEL P. DUFFEY, UNDER SECRETARY OF WAR FOR ACQUISITION AND SUSTAINMENT

The Honorable Michael Duffey—was sworn in as the Under Secretary of War for Acquisition and Sustainment (USW(A&S)) on June 5, 2025. In this role, he is responsible to the Secretary of War for all matters pertaining to acquisition; contract administration; logistics and materiel readiness; installations and environment; operational energy; nuclear deterrence, chemical, and biological defense; the acquisition workforce; and the defense industrial base.

Over the past two decades, Mr. Duffey has served in several senior roles throughout the Department of War, including Deputy Chief of Staff to the Secretary of War, Chief of Staff to the Under Secretary of War for Research and Engineering, as well as multiple positions in the former Office of the Under Secretary of Defense for Acquisition, Logistics and Technology. Throughout his tenure in the Department, Mr. Duffey has led efforts to improve the application of systems engineering in the Department's development and procurement of our most complex weapon systems, established a capability to prioritize and protect critical technologies from exfiltration by our adversaries, and led teams developing critical technologies that maintain U.S. technological advantage on the battlefield.

In addition to his leadership at the Department of War, Mr. Duffey served as the Program Associate Director for National Security at the Office of Management and Budget (OMB), where he oversaw the entire $1+ trillion national security budget for the United States Government, including the Departments of War, State, Veterans Affairs, the Office of the Director of National Intelligence, the Central Intelligence Agency, and several international economic agencies. In his role at OMB, Mr. Duffey worked with Department leadership to publish a robust 30-year shipbuilding strategy prioritizing U.S. maritime superiority, and led efforts to fund nuclear modernization and the standup of the U.S. Space Force.

Under Secretary of War for Acquisition and Sustainment Duffey is a graduate of the University of Wisconsin.

THIS PAGE INTENTIONALLY LEFT BLANK

PANEL 14. FORGING THE SHIELD: COLLABORATIVE DEVELOPMENT FOR IMMEDIATE WARFIGHTING ADVANTAGE

Thursday, May 7, 2026, 1015 – 1130 ET (0715 - 0830 PT)

Panel Summary:

While high-level policy sets the strategic intent, the actual delivery of overmatch occurs where tactical operators meet technical innovators. This plenary panel focuses on the "Innovation Operating System"—the execution-level organizations and their industry counterparts who are transforming requirement-setting from a linear process into a rapid, iterative loop. By bringing together the organizations responsible for scaling drones, AI, and hypersonics with the operators who field them and the companies that build them, this session explores how to move past prototypes to achieve true battlefield scale.

Chair: Raymond D. Jones, Col. (Ret.) USA, Professor of the Practice, Department of Acquisition, Finance and Manpower, Naval Postgraduate School

Panelists:

Military Operator – *RADM Michael Mattis, USN, Director, Strategic Effects, Commander U.S. Naval Forces Europe/Africa Commander, Task Force-SIX SIX (CTF-66)*

Traditional Defense Powerhouse: Top-Tier Defense Prime – *Chris Mang, Vice President, Strategy & Business Development, Lockheed Martin Rotary and Mission Systems (RMS)*

Military Operator – *Ariel Dvorjetski, PEO for Weapons, Israeli Air Force*

Non-Traditional Industry Disruptor – *Kyle Tucker-Davis, Director of Data Engine for Scale AI Public Sector*

Raymond D. Jones, Col. (Ret.) USA—retired as a Colonel from the U.S. Army in 2012 and is a Professor of Practice with the Graduate School for Business and Public Policy at the Naval Postgraduate School. His last assignment in the Army was as the Deputy Program Executive Officer for the Joint Tactical Radio System (JTRS).

Additionally, he served as the Military Deputy for the Director of Acquisition Resources and Analysis in the Office of the Under Secretary of Defense for Acquisition Technology and Logistics (USD(AT&L)), managed three Major Defense programs for the DoD in addition to his many operational and research and development assignments.

He graduated from the U.S. Naval Test Pilot School in 1995 and is 1983 graduate of the United States Military Academy. He has a Bachelor of Science degree in Aerospace Engineering, a Master of Science Degree in Aeronautical Engineering from the Naval Postgraduate School, a Master's in Business Administration from Regis University, a Master's Degree in National Resource Strategy from the Industrial College of the Armed Forces and is currently a PhD candidate with the Graduate School of Information Sciences at the Naval Postgraduate School in Monterey California.

RADM Michael Mattis, USN—is from Fullerton, CA, graduated with distinction from the United States Naval Academy, Annapolis, MD in 1994 and Oxford University, England in 1996 as a FitzGerald Scholar. At sea, Mattis qualified in submarines, as a nuclear engineer and as a ballistic missile officer while serving on USS HENRY M. JACKSON (SSBN 730)(GOLD). He also deployed aboard USS ALABAMA (SSBN 731)(GOLD) and USS COLUMBUS (SSN 762) during his Junior Officer tour from 1998-2000.

Ashore, Mattis served on the staff of the Chief of Naval Operations, Washington, DC as an Assistant for Political-Military Affairs in 1994 and at the Naval Satellite Operations Center (NAVSOC), Point Mugu, CA from 2000-2002 as Satellite Operations Department Head and later Operations Director (N3/5/6). He transitioned to the Navy Reserve in 2003.

Returning to Active Duty in 2008-2009, Mattis, trained in Army Civil Affairs and deployed to Baghdad, Iraq as part of Operation Iraqi Freedom (OIF) in support of the XVIII Airborne Corps/Multi-National Corps-Iraq (MNC-I. In 2020 Mattis returned to Active Duty supporting OPNAV N7 as the Space, Information Warfare and Integrated Fires Lead for the Tri-Service Maritime Strategy and led COVID-19 relief efforts in Los Angeles (LA) as the shore OIC and later Commander, Task Force-LA. He returned to Active Duty again in 2021-2022, serving as Director, Navy Space Command (NAVSPACE), leading the stand-up of NAVSPACE as the Navy Service Component to US Space Command (USSPACECOM). Mattis assumed his current roles as Director, Strategic Effects for Commander U.S. Naval Forces Europe/Africa and Commander, CTF-66 in October 2023.

Mattis' awards include the Legion of Merit (3), Bronze Star Medal, and Meritorious Service Medal (4). He is a graduate of USMC Command and Staff College and Joint Forces Staff College Joint Professional Military Education programs, National Defense University Reserve Component National Security Course, Naval War College Executive Level OLW Course and Naval Postgraduate School Navy Senior Leader Seminar. He is a Joint Qualified Officer and qualified as an Information Warfare (Space Cadre) Officer.

Chris Mang—is vice president of Strategy & Business Development for Lockheed Martin's Rotary and Mission Systems (RMS) business area, a $17.2 billion enterprise that employs approximately 35,000 people and manages more than 1,000 programs. He is responsible for leading RMS's global strategy and business-development activities, delivering innovative, competitive, and reliable solutions that enable Lockheed Martin customers to create a more secure and prosperous world. Chris also oversees strategic planning, investment, acquisitions, and new-business concepts for each of the four lines of business that comprise the RMS portfolio.

Mang previously served as the vice president of Strategy & Business Development for Lockheed Martin's Missiles and Fire Control (MFC) business area, where he developed a comprehensive growth strategy supporting operations in more than 50 countries across 50+ product and service lines. Earlier in his MFC tenure, he was vice president of Strategy, leading multiple teams in strategy development, cross-business-area collaboration, portfolio management, alignment of new-business resources, mergers and acquisitions, and government affairs.

Throughout his career, Mang has led several critical initiatives for Lockheed Martin, including serving as vice president of Strategy & Business Development for Air Dominance Weapon Systems (ADWS) and Tactical & Strike Missiles (TSM). In these roles he led multiple teams and was responsible for overseeing line of business growth, customer relationships, and partnerships in the US and in 30 countries worldwide. This portfolio included air-to-air missiles, ground and air launched weapons, anti-tank weapons, long range cruise missiles, hypersonic weapons development and special programs.

Mang joined MFC in 2012 from Lockheed Martin Government Affairs where he was responsible for business development in the Navy unmanned air vehicle, weapons, sensors, and classified program areas. He started with Lockheed Martin in 2003 at what was then Simulation, Training and Support after a 10-year career in the Navy as an E-2C Hawkeye Naval Flight Officer.

Mang graduated from the United States Naval Academy with a Bachelor of Science degree and has a Master of Business Administration from Old Dominion University.

Colonel Dr. Ariel Dvorjetski—Brigadier General (Selected) Ariel Dvorjetski is the Program Executive Officer for Weapons and Commander of the Armament and Electronic Warfare Department, Israel Air Force Materiel Directorate.

He is responsible for the planning and execution of all life cycle activities for air-delivered munitions for the Israel Air Force Weapons Portfolio. He also provides strategic and technical guidance to the workforce of military, civilian, and contractor personnel leading discovery, development, and integration of all warfighting weapons technologies for the Israel Air Force.

The Armament and Electronic Warfare Department is responsible for designing, developing, producing, integrating, fielding, and sustaining a family of air-to-ground and air-to-air munitions for the Israel Air Force to defeat a spectrum of enemy targets. The Department advances applied research for seekers, navigation and control, image processing, munitions integration, warheads, fuzing, explosives, and technology assessment methodologies, as part of its responsibility for managing the Air Force Weapons Portfolio Science and Technology program.

BG(s) Dvorjetski was nominated to assume the role of Head of the Materiel Directorate in July 2026.

Previously, he was the head of Aircraft Programs and Engineering Center, Israel Air Force Materiel Directorate from July 2020 to August 2023, where he was responsible for total life cycle management of all aircraft systems. he handled the acquisition and initial fielding of new aircraft systems and upgrades.

BG(s) Dvorjetski was the Air Force Technical Airworthiness Authority, the single independent official authorized to make determinations that air systems are safe for flight and approved for flight operations.

BG(s) Dvorjetski was born in Tel-Aviv in 1978. In 2000 he joined the Israel Air Force after graduating the faculty of Aerospace Engineering at the Technion (Israel Institute of Technology) with distinction, as part of the IDF Academic Reserve program. He also holds Ph.d degree in Aerospace Engineering from the Technion. His research thesis was focused in basic fundamentals of Counterflow Spray Diffusion Flames, which earned him the International Bernard Lewis Fellowship award at Montreal, 2008. BG(s) Dvorjetski was also recognized by the Israel National Wolf Foundation for academic excellence. Throughout his career he has published more than 20 papers in world-class peer-review journals.

BG(s) Dvorjetski is also a graduate of the Israel National Defense College, 2023, where he earned the degree of M.A. in Security and Strategic Studies.

Kyle Tucker-Davis—is currently the Director, Data Engine for Scale AI Public Sector, focused on the development of reliable AI systems for both Computer Vision & GenAI applications. Prior to his current role, Kyle spent multiple years driving the creation of AI training data for national security applications. Before Scale AI, Kyle spent 10 years as an Infantry Officer in the United States Marine Corps, where he served in various operational and instructor roles. Kyle's academic background includes a B.S. in Mathematics from the United States Naval Academy, M.A. in Military Studies & Leadership, and Graduate Certificate in Artificial Intelligence.Prior to the JAIC, Major Tadross, an Air Traffic Control Officer, was hand selected as a member of the Secretary of the Navy's Innovation Advisory Council, where he conducted independent research at the MITRE Corporation on integrating AI into Aviation Command and Control. His previous tours included Marine Air Tactical Command Squadron 18 in Okinawa, Japan where he served as the Senior Air Coordinator and supported exercises in South Korea, the Philippines, and Australia. He has a Bachelor of Science degree in Mechanical Engineering from Old Dominion University and a Masters in Technology Management from Georgetown University.

THIS PAGE INTENTIONALLY LEFT BLANK

PANEL 15. THE MODERN ARSENAL: SCALING PRODUCTION THROUGH HETERODOX ACQUISITION AND STRATEGIC INVESTMENT

Thursday, May 7, 2026, 1145 – 1300 ET (0845 - 1000 PT)

Panel Summary:

Accelerating warfighting capabilities requires a fundamental shift from traditional, often confrontational procurement models to integrated strategies that prioritize industrial scale and coalition readiness. This panel explores the "Large Lot Procurement" (LLP) typology designed to maximize munitions production at full economic rates, alongside a proposed framework for synchronizing U.S. munitions demand with allied co-production and equipping needs—specifically regarding regional stability in the Indo-Pacific. Furthermore, the panel examines the emerging role of the Department as a strategic investor, using direct equity in private firms to catalyze private capital and secure critical supply chains.

Chair: Dr. Dale L. Moore, Acquisition Sciences, Program Management Faculty, Naval Postgraduate School

Panel Presenters:

> **Large Lot Procurement for Munitions Acceleration: A Novel Acquisition Typology to Catalyze a Defense-Industrial Renaissance –** *Ryan Burgess & Darryl R. Chew, Office of the Undersecretary of War (Comptroller)*
>
> **A Strategic Investor: Federal Equity Investing Approaches to Advance National Security –** *Sam Moyer, National Defense Industrial Association (NDIA)*
>
> **Strengthening the Arsenal of Deterrence: Modernizing the DoD Munitions Requirements Process to Include Allied Equipping –** *Sara Eighmey, U.S. Indo-Pacific Command*

Dr. Dale L. Moore— is the Chair of the Maryland Aerospace and Technology Commission. He currently serves as the President of the Southern Maryland Navy Alliance and is an Intermittent Lecturer at the Navy Postgraduate School in the Acquisition Sciences Group of the Department of Acquisition, Finance, and Manpower. He chairs the Southern Maryland 2030 Workforce Board of Advisors and is a member of the Maryland Military Installations Council, the Southern Maryland Economic Development Association, the Patuxent Partnership Board of Directors and the St. Mary's County Economic Development Commission.

Dr. Moore is the Founder and President of The Moore Group LLC. He has 43 years of experience across the full life cycle of DoD Acquisition. He has successfully led enterprise-wide strategic operations and organizational change and transformation initiatives, is a regular public speaker and panelist as well as a published academic researcher.

Dr. Moore has served in the Department of the Navy (DON) as the Director, Strategy and Innovation, DASN(RDTE) where he led development of the Department of the Navy's (DON) 30 Year Research and Development Plan, and supported the development of the 2045 Future Fleet Design and Architecture (FFDA), the Office of the Secretary of Defense Science and Technology Vision, and was selected as a Navy Innovation Advisory Council (NIAC) Fellow where he led the DON Strategic Thinking Community of Interest and the "DON Leadership @ the Edge Series."

Prior positions include: Naval Air Systems Command (NAVAIR) Director, Strategic Initiative Coordination and Execution, Assistant to the Commander of the Naval Air Warfare Center – Aircraft Division for Strategic Operations, Chief Systems Engineer for PMA-268 Navy Unmanned Combat Air System (UCAS/X-47B), Deputy Corporate Deployment Champion for "NAVAIR AIRSpeed" Continuous Process

Improvement including leading the development and release of the DON-wide Continuous Process Improvement Solicitation, and NAVAIR's National Competency Leader for Materials Research and Engineering.

Dr. Moore holds an Ed.D. from The George Washington University; an MS with Distinction from the Naval Postgraduate School; and certificates of completion from the M.I.T Sloan School of Management and School of Engineering; and a BS in Mechanical Engineering from the University of Delaware.

Dr. Moore has received the Department of the Navy Meritorious Civilian Service Award three times.

Large Lot Procurement for Munitions Acceleration: A Novel Acquisition Typology to Catalyze a Defense-Industrial Renaissance

Ryan L. Burgess—is currently a Financial Management Analyst for the Office of the Secretary of War Comptroller/Deputy Chief Financial Officer (OUSW[C]/DCFO) Enterprise Financial Transformation (EFT) Directorate. Mr. Burgess provides programmatic support for the Joint Task Force – Audit (JTF-A), a Comptroller-wide initiative to accelerate Department progress at meeting the Congressionally mandated 2027 audit deadline.

Formerly a John S. McCain Strategic Defense Fellow, Mr. Burgess holds an MA in Global Security Studies, a Certificate in Intelligence from Johns Hopkins University, and a BA in Philosophy from the University of Chicago.

Darryl R. Chew—is currently a Senior Budget Analyst for the Office of the Secretary of War Comptroller/Chief Financial Officer (OUSW[C]) Program/Budget, Investment. Mr. Chew advises senior executive staffs, Defense Agencies, and Services on mission goals, auditable objectives, and funding strategies. Mr. Chew manages the Golden Dome for America initiative, the U.S. Missile Defeat and Defense (MDD) portfolio, and key portions of the munitions portfolio for the Department of War (DoW) Comptroller, where he implements the Secretary's and Administration priorities and ensures they are consistent with the Office of Management and Budget (OMB) and DoW policy, properly costed, and executable within the budget time frame. He provides the DoW Comptroller senior leadership with pertinent financial information and budget analysis to permit the senior leadership to make timely informed decisions regarding key defense acquisition programs.

Darryl Chew spent several years on Capitol Hill, in the House of Representatives and U.S. Senate, as a Legislative Assistant. He also was the Program and Legislative Liaison for the Missile Defense Agency before becoming its Comptroller Chief of Staff. He is Level III Business and Financial Management certified. He is a graduate of the Georgetown University School of Foreign Service with degrees in International Politics and National Security.

Abstract

Large Lot Procurement (hereinafter, Large Lot or LLP) was designed to address one of the primary structural constraints in munitions production: the current misalignment between government and private sector incentives. The proximate goal of this contracting typology is to incentivize vendors, at the prime level as well as lower tiers, to invest their own funds into facilities, personnel, and R&D. Accomplishment of this immediate aim, from production efficiencies as well as this additional investment, is then leveraged for additional manufacturing of end items. While, as originally conceived, Large Lot Procurement was intended to catalyze munitions production and supply chain development, it could easily be applied to other categories of defense articles that are like munitions.

This essay has eight sections:

1. The Need for Acquisition Transformation
2. Overview and Brief History of LLP
3. The American Arsenal of Democracy and Its Modern Discontents
4. Implementation Considerations
5. Recent Policy Decisions and Their Relevance
6. Modeling Why It Works
7. Large Lot Beyond Munitions
8. Lessons Learned: How to Spur an Acquisition Renaissance

Section I: The Need for Acquisition Transformation: A Joint Urgent Operational Need

If there is a single area of defense policy that has nearly unanimous consent across the political spectrum, it is the need for critical munitions. Analysts from news outlets, think tanks, industry, universities, and other research institutions have been raising this issue with remarkable frequency. Their collective message, all within the past three years, includes such headlines, in relative order of urgency, as:

1. "The U.S. Defense Industrial Base Is Not Prepared for a Possible Conflict with China" (Jones, 2023)
2. "Fleet at Risk as Navy Struggles to Refill Missile Stores" (Frazier, 2025)
3. "The U.S. Military Faces a Critical Missile Gap Against China" (Green, 2025)
4. "America Doesn't Have Enough Weapons for a Major Conflict" (Leonard, 2025)
5. "US Only Has 25% of All Patriot Missile Interceptors Needed for Pentagon's Military Plans" (Lowell, 2025)
6. "The Crumbling Foundation of America's Military" (Bowden, 2024)

These were all issued before the initiation of Operation Epic Fury in February 2026 which has resulted in the expenditure of massive quantities of munitions at rates that greatly exceed current production. If the news of the past three years lends credence to the general conclusion that current procurement practices are not enough for expected needs, the events in the Middle East have changed the game. Not only do we not have enough munitions for a potential future conflict, but we may also not have enough for the current one if it is protracted.

Conflicts within this timeframe—despite their relatively short duration and limited overall U.S. involvement—have finally laid bare the fact that our munitions production and procurement processes are fundamentally insufficient for our wartime operational requirements. A 2025 article in *The National Interest* put the editorializing headlines in stark perspective: "The [U.S.] military used one fourth of its Terminal High Altitude Air Defense (THAAD) interceptors to defend Israel [for 12 days]. Imagine how many it would need to defend Taiwan" (Eaglen & Harrison, 2025). The same incongruence does not just pertain to advanced capability munitions, featuring exquisite components, complex production processes and supply chains, such as PAC-3 Missile Segment Enhancement (MSE) and THAAD's interceptors. To name just one example, "[a] single Ukrainian [artillery] battery can fire more 155mm rounds *in a day* [emphasis added] than some American units used during the [1991] Iraq War" (Soliman, 2025). Soliman's diagnosis is blunt.

> America's military is built for the wrong kind of war. The United States has optimized its defense industry for short, high-tech conflicts using precision weapons, but current global wars require sustained, large-scale production of conventional munitions. … Recent wars demonstrate that even the most advanced weapons are useless without the techno-industrial capacity to produce them in massive quantities, … as seen with Germany's defeat in World War II due to an inability to match Allied production. (Soliman, 2025)

While he couches his criticism by noting that the allure of "the cult of technology and the precision weapon" was "seductive and, within its narrow parameters, correct," the conflicts the United States engages in today require exponentially greater quantities of munitions and materiel than originally planned or budgeted for (Soliman, 2025). This brutal logic has even worse consequences the further right the timeline goes: The U.S. would confront even more challenging scenarios in any large-scale conflict with the People's Republic of China (PRC). If the Navy, for example, "burned through 30 years of missiles in [only] 15 months" when facing a non-peer adversary in a limited conflict of short duration, the magnitude of the production disparity is catastrophic in the face of a peer or near-peer (Frazier, 2025). In other words, we are out of time. We need several orders of magnitude more munitions than we currently produce, and we need them immediately.

The U.S. defense community therefore faces the following challenges:

1. Current quantities and production levels are severely and consistently below requirements.
2. Past budget cycles and acquisition reform efforts have failed to materially close the persistent gaps between production and operational expenditure.
3. Recent investments in the defense and organic industrial bases will not yield the required improvements in even the medium-term time horizon.
4. The United States has limited fiscal breathing room, and today's permissive budgetary environment may not persist beyond the next several years.
5. Our principal adversaries already possess much higher production capacity and are keenly aware of this disparity.

Interlocutors' and analysts' proposed solutions generally seek to leverage existing U.S. strengths and competencies. Eaglen and Harrison (2025), in *The National Interest*, argue that "[t]he most promising solution for the fastest return is to leverage commercial innovation." Wasser and Sheers (2025), from the Center for a New American Security (CNAS), argue for more creative and flexible contracting and funding strategies. In a Heritage Foundation report entitled "A Strategy to Revitalize the Defense Industrial Base for the 21st Century," Greenway et al. (2025) use almost the same verbiage.

Further commentaries from Brookings and the Assistant Secretary of the Army for Acquisitions, Logistics, and Technology (ASA[ALT]), to name just two examples, suggest the need to put the Defense Industrial Base (DIB) on a "wartime footing," to quote the Secretary (McGinn & Cook, 2025). The government should not only be "[d]irectly incentivizing or even funding companies to develop spare production capacity and thus accommodate increased demand" (O'Hanlon & Rocha, 2024). The Department must provide "a coherent national demand signal" as a baseline requirement for achieving the necessary wartime footing (McCoy, 2025). Large Lot Procurement will provide the consistent, coherent demand signal that is required for mission-critical industry-government collaboration.

Section II: An Overview and Brief History of Large Lot Procurement

Large Lot Procurement (LLP) predates the munitions and procurement reform *zeitgeist* by over a decade. The progenitor of this idea, a past Director of OUSW Comptroller Program/Budget's Investment Directorate (INV), began to develop the idea in the early 2010s. His deep professional experience throughout the Planning, Programming, Budgeting, and Execution (PPBE) process, in tandem with his contacts throughout government and the defense industry, broadly informed his insightful understanding of the fundamental challenges facing the Department's efforts to translate its procurement appropriations into deliveries on time and at budget.

Most importantly, Mr. Roberto Rodriguez possessed a clear-eyed understanding of the direct causes of the Defense Industrial Base's (DIB's) key difficulties that the Department must now urgently address. He leveraged his deep historical and organizational knowledge to address what he correctly ascertained to be one of the primary structural reasons why munitions procurement so often results in delayed deliveries, over-budget programs, and missiles that perform below expectations. Namely, Mr. Rodriguez created a framework that is designed to incentivize contractors to invest their own resources and increase production while assuring that the government fulfils many of the requests the private sector has been making that would, in their view, suffice as a consistent demand signal. In other words, LLP aligns the incentives of the private and public sectors to produce a non-zero-sum game-theoretic landscape for both primary actors.

Large Lot coalesced between 2014 and 2015. Mr. Rodriguez's rationale for developing an evolution of existing MYP contracting and financing tools stemmed from two primary realizations. Firstly, contractors lack—in their view—sufficient financial incentives and demand signals to justify additional investment of their own capital into facilities, tooling, long-lead items, etc.[1] Second, the government, for its own part, believes contractors are shirking their duty to the government by not delivering on time and on budget. This claim notwithstanding, the government may not understand precisely why private sector companies with strong financial and capital markets access fail to translate department investment and procurement orders into industrial capacity to support mass production.

Large Lot Procurement functions by executing multiple manufacturing programs under a concurrent, multiyear procurement strategy (Rodriguez, 2023). It enables synergies across different but related programs that drive greater production capacity, accelerate deliveries, and drive down unit costs. LLP is a stepwise evolution of existing MYP contracting. Critically, Large Lot Procurement expands the use of incremental financing, the savings derived from which are applied toward additional end item deliveries and higher production and procurement rates. This is referred to as "Buy-to-Budget" financing. Crucially, the financial windfall from production efficiencies is reinvested into production as opposed to being returned to government coffers.

As initially envisaged, implementation of this contracting typology included two primary lines of effort, each aligned to one primary contractor:

1. Lockheed Martin – The final assembly of the AGM-158B Joint Air-to-Surface Standoff Missile – Extended Range (JASSM-ER) and AGM-158C Long Range Anti-Ship Missile (LRASM) are at the same facility in Troy, AL, and both processes employ a common set of subcontractors, production tooling and machinery, and raw materials and intermediary components.
2. Raytheon – The Standard Missile-6 (SM-6) and AIM-120 Advanced Medium-Range Air-to-Air Missile (AMRAAM) are assembled in the same facility in Tucson, AZ, and employ a common set of subcontractors, production tooling and machinery, and raw materials and intermediary components.

While MYP supports the health of the Defense Industrial Base by creating program stability—critically important for lower-tier subcontractors and Small and Medium-sized Enterprises (SME)—Large Lot goes further than traditional MYP by proactively addressing known bottlenecks and inefficiencies through upfront industrial base financing. It therefore builds upon one of the strengths of MYP—financing increased production at common manufacturing lines—which allows Original Equipment Manufacturers (OEM) and subcontractors to plan for and synchronize production in an orderly, more efficient fashion. This, in turn, mitigates against fitful,

[1] Furthermore, the government and private industry have different motives. While contractors act in service of national security, they are for-profit entities that have a duty to shareholders. The government does not have a profit motive, and its mission overrides all other considerations. This misalignment in ends is another reason why this issue has resisted previous efforts at reform, like 2009's Weapon System Acquisition Reform Act.

“stop-start” production schedules. Unfortunately, common in defense production, these factors introduce manufacturing inefficiencies, complicate logistics of raw materials and critical subcomponents, and reduce labor efficiency. As a direct result of the received acquisition process, “stop-start” production is particularly harmful for the manufacturing and assembly of the most advanced munitions owning to their greater internal complexity and inclusion of exquisite components. Large Lot is therefore best suited for accelerating the production and delivery of the most advanced munitions.

Finally, LLP differs from MYP in its government-side formulation:

1. The collection of concurrent production programs is managed at the Program Executive Officer (PEO) level, as opposed to the lower-echelon management of standard MYP programs at the Program Manager (PM) level.[2]
2. LLP features vertical integration (parallel workflows across the various weapon systems) as well as horizontal integration (executing one program across multiple fiscal years). MYP only includes the latter.
3. All annual procurement quantities are fixed at the Economic Production Rate (EPR) at a minimum, if not the full existing, advertised production capacity. In the originally envisaged program, annual procurement quantity is increased from Minimum Sustaining Rate (MSR) to EPR.
4. Inclusion of an upfront industrial base financing package to invest in production line efficiency and crucial capital projects. MYP as conventionally understood does not include direct public investment in industrial base needs.
5. Economic Order Quantity (EOQ): LLP also includes financing for long-lead item procurement. This is primarily targeted toward lower-tier subcontractors that produce intermediary components, basic components, or raw materials. MYP, in contrast, is oriented toward final assembly by prime contractors.
6. Buy-to-Budget, wherein savings generated by EOQ are used to procure additional end items. Savings from MYP, on the other hand, are not applied for further production quantity.
7. Concurrent execution: Work is executed in parallel (multiple systems over fiscal years) and the same place (OEM/factory/subcontractors).
8. Overlapping: Two or more MYP for the same end item. LLP is geared toward programs that produce advanced All Up Rounds (AUR); the benefits primarily accrue from rationalizing the production processes involved in assembling complex munitions.

[2] The Administration’s efforts to establish Direct Report Program Managers (DPRMs) and Program Acquisition Executives (PAEs) to manage large acquisition programs may result in a different organizational structure than LLP originally contemplated.

Table 1. Large Lot Procurement versus Multi-Year Procurement

Multi-Year Procurement	Large Lot Procurement
Similarities	
Minimum savings requirement Stable requirements Stable funding/availability Realistic cost estimate	
Differences	
Single contract	Parallel, overlapping and concurrent MYP contracts
Savings recouped	Savings reinvested into weapon systems
Design Stability	Flexible design
Unique production facility	Common production facility with another weapon system
Unique manufacturing process	Common manufacturing processes
MYP allows for reduction of long lead items	EOQ funding can be applied toward reducing long lead-time constraints

The primary benefit is in savings/cost avoidance. As originally planned, implementation of LLP provides a 10% annual APUC saving during the contract period (Rodriguez, 2023). Furthermore, LLP preserves and improves industrial base capacity and efficiency at both prime contractors and subcontractors. It also stabilizes the defense and organic industrial bases and incentivizes private sector investment in facilities, tooling, and production capacity; leverages existing industrial capacity to control weapon system costs; and uses cost savings derived to procure additional weapons (Rodriguez, 2010b).

Section III: The American Arsenal of Democracy and Its Modern Discontents

Summarizing, let alone synthesizing, modern commentary—or even the brief historical accounts—on U.S. industrial mobilization is beyond the scope of this essay. The revitalization of the U.S. Defense Industrial Base before and during the Second World War, the importance of its success to the U.S and the Allied war effort, and the lessons from this endeavor for the present are all important subjects. Reindustrialization, particularly with respect to the defense and organic industrial bases, is a well-deserved *cause célèbre* across the political spectrum.

Freedom's Forge: How American Business Produced Victory in World War II, a 2013 book by Arthur Herman, is particularly notable in the recent scholarship on this topic. It tells the well-trodden story of U.S. mobilization from the perspective of two important American industrial figures, William S. Knudsen and Henry J. Kaiser. Herman highlights several key factors that were essential to developing the U.S. industrial might that was a significant factor in ensuring Allied victory (Culclasure, 2013). Herman is particularly insightful in laying out the basic principles that allowed the U.S. to successfully mobilize its industry, especially given the diminutive military and low level of government-industry coordination in the interwar period. Per *Freedom's Forge*, the primary precondition for this ultimate success was a fundamental, clearly defined division of labor between the public and private sectors:

1. The government decided what would be produced, pair for it, and ensured the availability of basic materials needed (Smith, 2022).
2. Industry, on the other hand, undertook the actual production, from design to assembly, supply chain management, and labor policy (Smith, 2022).

Notwithstanding regular—and often serious—disputes between the government and its private-sector partners, to say nothing of labor strikes, supply chain issues, and other unforeseen challenges stemming from building a wartime industrial machine from the ground up, Herman frames the result in plain terms. By the conflict's end, "the U.S. managed to out-produce the rest of the Allies combined, while devoting a much smaller portion of its economy to the military than other nations and even *increasing* civilian consumption" (Smith, 2022).

The book's insightful analysis, its framing of a subject of national importance with an elegant model of political economy, and its straightforward lessons for contemporary challenges have all been well reviewed by the national security commentariat. To wit, to respond to today's challenges in the Defense Industrial Base, we need to restore a healthy, mutually beneficial relationship between the private sector and its ultimate master. Large Lot Procurement is well-suited as a critical tool to achieve this result.

Section IV: Implementation Considerations

Probable hurdles to effective execution include, but are not limited to:

Budgetary Instability. The single greatest obstacle to enduring effectiveness of LLP (or any other procurement reform) is the Congressional appropriations process. The cyclical, partisan-forward nature of budgeting in the legislative branch—to say nothing of the unstable, often erratic, schedule of the release of the yearly President's Budget Request—makes it politically and structurally difficult for the Department to *credibly* commit funds, at the scale required for efficacy, for a multiyear procurement cycle spanning multiple legislative sessions. Absent changes to the appropriations process itself—which is beyond the scope of this essay and is even less likely than the reforms contemplated here—a future Congress, or presidential administration, could unilaterally alter priorities or funding. This would leave both the Department of War and its industrial partners, all of whom have proffered funds, reputational credibility, and institutional inertia on a new process, in a precarious legal and financial position.

Internal negotiations in 2022 highlight organizational tensions in executing long-term munitions production programs. Specifically, SASC's preferred proposal for weapons systems multiyear procurement was not true MYP, insofar as it failed to include EOQ financing and thereby lacked one of the key attributes that produce LLP's benefits. Moreover, it did not create an obligation on future Congress and Department to complete the planned procurement without invoking termination liability charges (Rodriguez, "Munitions Multiyear Proposal in Draft NDAA 2," 2022).[3] Removing Department financial obligation will *very likely* reduce the potential for mass industry buy-in. All parties need to have credible stakes in the execution of the program for its duration, or the structural reasons for its success—solving the information asymmetry problem to allow for effective coordination—will not hold.

Technological Obsolescence. LLP requires that the Department and its vendors make a significant investment in particular munitions, each with specially sourced components. Committing to a five- or more -year procurement of a particular weapons system is inherently risky in a world characterized by rapid, accelerating technological change. As originally envisaged, weapon configurations would largely be 'fixed' at the time the contract is finalized (Rodriguez, "Large Lot Procurement – DepSecDef – November 2010"). However, funding stability may be amenable to spiral capability additions with government approval (Rodriguez, "Large Lot Procurement – 3 STAR Briefing – DRAFT2," 2010). As a result, Large Lot is particularly useful for programs that allow for improvements to be made over time in the form of

[3] Of the sort required by, for example, DFAS 252.239-7007 or 52.241-10. Asking contractors to commit resources absent either the expected payoff (from LLP being executed) or insurance in the form of cancellation liability runs against their fiduciary obligations.

improved variants. Current systems that may be particularly amenable to Large Lot include PrSM, LRASM, or even less exquisite systems like LUCAS.

Industrial Base Fragility. The LLP concept assumes, as a matter of course, that the defense and organic industrial bases already possess the latent capacity and willingness to scale up if a sufficient government demand signal is provided. This is not a given. Decades-long industry consolidation, facilities offshoring, and an aging skilled workforce, to name just a few issues, have left multiple supply chains brittle and dependent upon a small number of suppliers and diminishing manufacturing sources. Even with a contract that guarantees payoff over many years, there is an inherent concern that the industrial base may simply not be up to the task of sourcing adequate machine tools, recruiting and maintaining skilled labor, long lead-time materials, and testing and evaluation processes to expand production at the rate that is contractually required. This uncertainty demands comprehensive market research into the state of the industrial base before the contract is signed.

First-of-its Kind (FOAK) Risks. A final, critical impediment is that the evidence presented for the virtues of this proposal is theoretical and notional. Speculative reasoning is used to prove the points explicated here. Acquisition executive leadership will be taking significant professional risks by implementing a first-of-its-kind approach to acquisition that lacks prior actualization examples. In other words, the first *true* test of LLP would have no strict precedent. This issue is, in general, common to all FOAK projects, whether they are industrial demonstration plants, manufacturing processes, or conceptual (legal) innovations. Current Department initiatives, including but not limited to the Munitions Acceleration Council (MAC), ..., advancement of 7-year contract terms for MYP versus the standard five years are promising steps that begin to address FOAK risks and challenges. The results are not apparent yet and will not be for some time; but the efforts currently ongoing are a significant departure from the norm. Addressing a culture of risk avoidance may require upstream organizational and cultural changes within the acquisition workforce. For example, advance socialization of the idea and a binding agreement with CAPE to consider LLP throughout future Program Budget Review cycles may help abate institutional hesitance to FOAK ideas like LLP. The establishment of Program Acquisition Executives (PAE) is precisely the sort of organizational change that is required to enable changes in contracting practices like LLP.

Successfully executing Large Lot Procurement requires program managers to address the following considerations:

Funding Strategy and Considerations (Rodriguez, 2010a):

1. LLP funds 2+ years of production in one fiscal year.
2. USW(P) approves deviating from the 12-month funded delivery period policy.
3. To avoid “new start” requirements, particularly the validation process, and to avoid requalification of hardware, procurement quantities are retained in all covered fiscal years.
4. OUSW(C) works with OMB to establish new, unique outlay (expenditure) rates for LLP.
5. Congress would have to agree to the LLP acquisition strategy (as with MYP).
    a. The Department should pursue express approval for Large Lot Procurement in all forthcoming bills. The current rate of expenditure of critical munitions in the Middle East means that getting more munitions into the hands of warfighters, as soon as possible, is a top national security priority.

    b. While Congress has periodically shown a willingness to adjust MYP requirements, the proposal must be announced in advance and socialized with relevant legislative stakeholders. Failure to get approval before legislative action risks the proposal getting delayed or watered down.

6. DoW assumes operational risk for the delay of new munitions.
7. No new tooling is financed by the government.
8. Objective is for the contract to be ROI neutral.
9. As noted above, Configuration and Engineering Change Proposals (ECP) would be restricted and would require USW(A&S) and USW(C) approval prior to implementation (Rodriguez, 2010a).
10. "Reclassification of Capital Budget / Account program."
11. Assumes a Lean 6 environment to maximize assembly rate/activity.
12. Production Strategy (Rodriguez, 2010a):
13. Prime and subcontractors will need to effectively phase work at their own facilities.
14. Contractors should adopt best practices, such as lean manufacturing principles.
15. Delivery schedule is negotiated prior to the signing of the contract, but it should be phased to the most efficient cost margin.
16. Contracting Strategy (Rodriguez, 2010a):
17. Omnibus contract vehicle for all production products.
18. Multiyear procurement concept from inception to end item delivery.
19. Firm Fixed Price (FFP) contract.
20. Priced annual options (for additional buys) by missile system.
    a. These can be rolled into the omnibus contract at government's discretion.
21. Negotiated price, fully indexed (at OMB rates) to cover inflationary costs.
22. Production risk cost share set to 50/50.
23. Earned Value Management System (EVMS).
24. Government option to use cost savings for additional purchases or for other requirements.
25. Payments set to 50% upon contract award, 25% based on assembly progress, and 25% upon delivery.[4]

Section VI: Modeling *Why* It Works

The issue at hand can be modeled as a coordination problem: The principal parties (i.e., the government and industrial firms) may earn a higher payoff if they successfully cooperate. This involves both entities coming to mutually beneficial agreements to ensure that the government receives the maximum quantity on time and at budget, while the private sector gets

[4] It is worth further studying the contract award spread to determine which best incentivizes contractor performance. Denominating the exact criteria for "assembly progress" beforehand and communicating the milestones to vendors will maximize chances of success by addressing the information asymmetry that characterizes government-contractor relations.

a more reliable order book into the medium-term future. While the results of positive coordination are superior to the status quo, it is not guaranteed that both actors will successfully cooperate. In fact, failing to do so (the so-called ‘defection’ option) still produces a positive, but not ideal, result. This often characterizes the landscape of munitions contracting and production. The government receives its deliveries—often late and over budget—and contractors’ profit, but the relationship between the two parties is often acrimonious. Both sides have a strong incentive to cooperate, which would produce mutual benefit, but structural factors in the relationship make it easier to stick with the (still beneficial, but suboptimal) status quo.

As such, the problem is a type of *coordination game*. Put in game-theoretic parlance, it is an assurance game colloquially known as the Stag Hunt. As originally posed, the two parties (hunters) can either jointly pursue a stag or individually hunt rabbits. The former is challenging and requires cooperation, but it carries a higher potential payoff. Importantly, defection still results in a positive, but less than ideal, outcome for any of the hunters in the party:

> In the Stag Hunt, what is rational for one player [the government] to choose depends on his beliefs about what the other [private actors] will choose. Both stag hunting and hare hunting are *equilibria*. That is just to say that it is best to hunt stag if the other player hunts stag and it is best to hunt hare if the other player hunts hare. … A player who chooses to hunt hare runs no rush, since his payoff does not depend on the choice of action of the other player, but he foregoes the potential payoff of a successful stag hunt. Here rational players are pulled in one direction by considerations of mutual benefit and in the other by considerations of [institutional] risk. (Skyrms, 2001)

The game-theoretic dynamics can be modeled with a simple Decision Matrix:

	Player 1: Hunt Stag	Player 1: Hunt Hare
Player 2: Hunt Stag	(10, 10)	(0, 7)
Player 2: Hunt Hare	(7, 0)	(7, 7)

In this matrix:

1. (10, 10): Both players cooperate to hunt the stag, resulting in the highest payoff for both.
2. (0, 7) or (7, 0): One player hunts the stag (unsuccessfully), whereas the other hunts the hare (successfully). This is less risky for both players, as the payoff variance over the other’s strategy is lower.
3. (7, 7): Both players hunt hares on their own, receiving a moderate (but still non-zero) payoff.

This dynamic neatly maps onto the government-industry relationship with respect to munitions production.

Munitions Decision Matrix:

	Department: LLP	Department: Business As Usual
Contractor: LLP	(10, 10)	(0, 7)
Contractor: Business As Usual	(7, 0)	(7, 7)

In this more illustrative matrix:

1. (10, 10): The Department and the Contractor cooperate to produce and deploy more munitions than otherwise. Highest payoff for both.
2. (0, 7) or (7, 0): One player attempts acquisition and contracting reform without involving the other. On the private side, this entails R&D and production investment absent binding signals from the government for future demand. Little or no payoff beyond business as usual for either player.
3. (7, 7): Business as usual. Received issues with the current system remain standing.

As applied to the munitions production situation, the stag hunt possesses two equilibria, with both players preferring one to the other:

1. Cooperative Equilibrium: The course of action with the highest payoff for both players requires trust and genuine cooperation. This is the so-called payoff-dominant strategy.
2. Non-Cooperative Equilibrium: Business as usual results on a lower reward, but it does not require cooperation. This is the so-called risk-averse strategy.

In coordination games like the Stag Hunt, players require knowledge about their counterparts: what they value, what they will do in certain circumstances, and the like. Absent this knowledge, *defections* that lead to realization of non-cooperative equilibria—less than ideal outcomes—will become the rule.[5] While this is a sub-optimal outcome for both parties, it is not the worst-case scenario. This game-theoretic typology therefore serves to illustrate the difficulties of multi *n* coordination by agents who (a) would both benefit from information sharing but (b) may also receive a discounted but still-positive payoff by way of defection from the cooperative equilibrium. The general dynamics of the assurance game therefore map isomorphically onto the munitions production and procurement situation:

Choosing the course of action that results in the cooperative equilibrium—and which produces the highest mutual payoff—*de facto* requires coordination and trust between players. This is seen in the current impasse between the government and its vendors:

1. The government often states that the private sector ought to invest more of its resources into R&D, facilitization, labor, etc. to meet production and fielding goals set by the NDS (e.g., TMR, Global Floor).
2. The government is also dissatisfied with what it views as subpar performance by the private sector: delays, price hikes, underperformance of delivered systems, etc.
3. In response, private industry claims that it lacks a consistent demand signal from the government that would grant it, and its stockholders, financial certainty that its investment will be recouped by way of future (and, ideally, increased) orders.

[5] The prototypical formulation of the Stag Hunt assumes that each player is a unitary entity with a single set of attributes. In contrast, the Department is a multifaceted entity, the components of which may have a range of positions on certain issues. Department entities that support the status quo are considered defections. Furthermore, defection by a subcomponent (e.g., a particular entity within the Department) may be enough to cause the entire Department to defect. It is therefore of the upmost importance to ensure top-down conformity once it is determined that Large Lot Procurement is the recommended course of action. A Direct Reporting Portfolio Manager approach may be precisely what is needed.

4. Industry also argues that balancing its fiduciary duty to shareholders requires that its non-mandatory expenditure on facilities, R&D, etc. has a relatively high likelihood of positive return.

LLP addresses the impasse between government and industry because it solves the "knowledge problem" that generally stymies the relationship between two players in assurance games. First posed by F. A. Hayek in 1945, this dilemma has two main components:

1. The *complexity* knowledge problem*:* The difficulty of coordinating individual plans and choices in the ubiquitous and unavoidable presence of dispersed, private, subjective knowledge. Trade secrets and internally privileged corporate operations prevent the government from possessing a so-called perfect understanding of the market. Government-mandated classification criteria and CUI dissemination controls (e.g., FED ONLY and NOCON) limit vendor understanding of government deliberations.[6]
2. The *contextual* knowledge problem*:* The epistemic fact that some knowledge relevant to coordination does not exist outside of the market context. Such knowledge is either created in the process of market interaction, tacit knowledge that is not consciously known …, or inarticulate knowledge that is difficult to express or aggregate (Kiesling, 2015).

Namely, "Hayek characterized the fundamental economic problem … as the coordination of actions and plans among dispersed agents with diffuse private knowledge" (Kiesling, 2015). Furthermore, prices serve a fundamental coordinating role in imperfect markets. Hayek adduced that "decisions facilitated by prices are indispensable in enabling … coordination to occur by providing a way to access dispersed, private knowledge" in a market characterized by compartmentalized information access (Kiesling, 2015). Prices are thus "knowledge surrogates" that "communicate the consequences of the realized actions and interactions" of the agents that participate in the market (Kiesling, 2015).

As defined above, Large Lot's implementation solves the knowledge problem by facilitating visibility, for all market participants, of key surrogates for otherwise compartmentalized (de facto secret) knowledge. The government's legal commitment to spending levels over the lifetime of the contract is a sufficient knowledge surrogate, to vendors, of the government's intentions. It is an advanced commitment to financing that addresses the private sector's competing demands to its customers and its shareholders. The inclusion of additional orders beyond the standard year-over-year budgeting practice therefore provides firms with an expected payoff that further incentivizes them to meet the government's performance requests.

Section V: Recent Policy Decisions and Their Relevance

Recent news articles and developments indicate that primary components of the original LLP proposal are being implemented today. For example:

1. On 23 February 2026, *Military & Aerospace Electronics* published an article analyzing a USAF task order to Lockheed Martin dated 13 February 2026. The analysis specifically mentioned Large Lot Procurement: "Large Lot Procurement refers to buying missiles in very high quantities under a consolidated contract instead

[6] In game theory, "perfect information" describes the situation in which all players have knowledge of all relevant system information. Also known as "no hidden information," the fact that this does not pertain in government-vendor negotiations means that the situation features "imperfect information." The knowledge problem between players inherently results.

of issuing smaller annual production awards to increase production capacity (thousands of missiles over several years); reduce cost per missile through economies of scale; stabilize the industrial base and supply chain; and accelerate fielding to meet U.S. and allied demand" (Keller, 2026).

2. On 17 February 2026, L3Harris Technologies was awarded a $400 million contract to produce SRMs and other control and guidance systems for THAAD (Miller, 2026). It is unclear if this contract is multiyear, but its focus on components below the AUR integration process speaks to the Department's awareness of supply and manufacturing constraints below the prime contractor level.

3. On 06 January 2026, Lockheed Martin (LM) announced a similar deal to triple production of the MIM-104F PAC-3 MSE. The firm's press release includes more detail than that provided by Raytheon. It notes that both the Department and Lockheed "will participate in the cost savings opportunity enabled by long-term demand certainty" for the munition (Lockheed Martin Corporation, 2026). The company's announcement also states that the joint DoW/LM framework:

   > [I]ntroduces a new model that provides long-term demand certainty, enabling industry investment, increasing production rates and driving operational efficiencies. It incorporates a collaborative financing approach designed to preserve initial cash neutrality, allowing industry to invest confidently to meet required production levels (Lockheed Martin Corporation, 2026).

1. The Department's press release on the DoW/LMT deal states that:

   > The framework agreement establishes the basis for negotiating a seven-year supply contract, subject to Congressional authorization and appropriations …. This facilitization strategy will be applied to multiple munitions procurement contracts over the next year …. … As part of the framework agreement, the DoW will work with key suppliers of PAC-3 MSE to deliver seven-year subcontracts to ensure facilitization investments and the production capacity of components also expand to meet the increased demand for all-up-rounds (U.S. Department of War, 2026).

2. The 06 January 2026 agreement between the Department of War and Lockheed Martin builds off the MYP contract, signed on 03 September 2025, to produce 1,970 PAC-3 MSE interceptors between FY 2024 and FY 2026 (Judson J. , US Army awards Lockheed record $9.8 billion missile contract, 2025).

3. Jason Reynolds, Lockheed's Vice President for Integrated Air and Missile Defense, commented earlier in the year that the firm looking at "efficiencies and streamlining and doing everything we can to stretch those [procurement] dollars to take [it] to a higher capacity, upwards of around 750 per year by 2027" (Judson, 2025b).

4. On 04 January 2026, Raytheon (RTX) and the Department agreed to increase production of six critical munitions (RTX, 2026): RGM/UGM-109E Block V Tomahawk Land Attack Missile (TLAM), RGM/UGM-109E Block Va Maritime Strike Tomahawk (MST), AIM-120 Advanced Medium-Range Air-to-Air Missile (AMRAAM; presumably the AIM-120-D3 variant), RIM-174 Standard Missile 6 (SM-6; unclear which variant[s]

are included), RIM-161C Standard Missile 3 (SM-3) Block IB, and RIM-161D Standard Missile 3 (SM-3) Block IIA.

5. Per the company's press release, the deal includes "up-to-seven-year agreements to … build on the company's previous investments to expand production" (RTX, 2026). Furthermore, production will take place at company facilities where multiple of the covered munitions are assembled (Tucson, AZ; Huntsville, AL; and Andover, MA). As noted above, Tucson and Huntsville are where many of these munitions are assembled; Andover is where many of the subcomponents are produced. On 15 January, Raytheon disclosed that it has invested over $115 million to secure a nearly 67% capacity increase at Huntsville, AL, for final integration of all Standard Missile variants (Judson 2025a).
6. On 31 July 2025, the Department awarded several contracts worth a total of $7.8 billion for many of the munitions envisaged for LLP. Commentators on the deal stated that these agreements marked "a major step in the Pentagon's new strategy of buying munitions in multiyear batches, a move officials say is necessary to provide a strong demand signal to industry and encourage companies to increase production capacity" (Hadley & Tirpak, 2025).
    a. Lockheed Martin Missiles and Fire Control was awarded $4.3 billion for five lots of JASSM as well as for lots' worth of LRASM. The contract announcement explicitly states that the "definitization modification (PZ0007) … [covers] large lot procurement for JASSM Lots 22-26 and LRASM Lots 9-12" (U.S. Department of War, 2025).
    b. Raytheon was awarded $3.5 billion for AMRAAM Lots 39-40. Besides production of AURs, the contract provides for production of missile components including telemetry systems as well as initial and field spares (Fletcher, 2025). Work will be completed in Q3 of FY 2031 (U.S. Department of War, 2025).
7. On 30 September 2024, the Department awarded LMT a $3.2 billion multiyear contract to procure AGM-158C LRASM and AGM-158B JASSM AURs until July 2032 (Losey, 2024).

Other private entities have also announced programs and projects aimed at meeting lower-tier supply and capacity gaps or entirely separate production lines:

1. On 25 September 2025, Avio USA (a subsidiary of Avio S.p.A.) signed a $26 million MoU with Raytheon in the process of becoming a second supplier of Mk 104 dual-thrust solid rocket motors (Raytheon, 2025). This follows a July 2024 contract for preliminary engineering and a subsequent November 2025 purchase order to fund the project through Critical Design Review (Raytheon, 2024, 2025). Subsequently, Avio announced that it would be developing a $500 million production facility in Virginia (Avio USA, 2025). This will build upon its plans to also assemble the Mk 104 at an existing plant in Colleferro, Italy (Kington, 2025).
2. On 17 September 2024, Kongsberg Defense & Aerospace announced that it will construct a missile production and maintenance facility in James City County, Virginia (Judson, 2024). The $100 million facility will help the company meet the existing multiyear demand for its RGM-184A Naval Strike Missile (NSM) and follow-on Joint Strike Missile (JSM) systems (Kongsberg Defense and Aerospace, Inc., 2026).

Finally, the Munitions Acceleration Council (MAC) has commendably acted with speed to greatly expand munitions production. However, the Department lacks a dedicated entity and process that is empowered to analyze whether this contracting approach is generating the contemplated gains in production volume and speed. The absence of an accountable entity or group should be deliberately addressed to avoid what may be repeated or common practice to quickly establish contract parameters outside of the standard process designed specifically to identify and prevent waste, fraud, or abuse.

Section VII: Large Lot Beyond Munitions

During the original idea development phase, the following demand-side attributes are deemed compatible with LLP (Rodriguez, 2010b): (1) Funded production rates, (2) firm military requirements, and (3) flexible production manufacturing facilities. Moreover, while the program was originally intended for munitions production, it is not inherently limited to missiles or other munitions. In other words, LLP has significant applicability beyond this type of defense articles. Ideal candidates for inclusion in an LLP scheme will have: (1) shared intermediary components, (2) diminished sources/supply chain challenges, (3) repeatable/mass producible end items, and (4) assembly of multiple systems at shared locations and/or contractor facilities. Potential areas of extension include, but are certainly not limited to, small satellites, radars and sensors, and turbines and rocket engines

Each of these meets the required attributes that are shared with the original, notional category envisaged for LLP. Satellites, particularly current and next generation small/microsatellites, are made on assembly lines; multiple variants are often produced by manufacturers at a single facility that brings together components from subcontractors; critical components must be produced *en masse* and often have significant long lead-time constraints and/or supply chain constraints. Spacecraft, like those the Department is looking to procure in the future, deviating from the traditional choice of exquisite, bespoke end items, makes small satellites an ideal candidate for an LLP-type contracting vehicle. Phased array radars, EO/IR sensor arrays, and aeroengines, particularly turbine engines, also meet these criteria. The Large Lot contracting typology possesses immense inherent promise beyond missiles and munitions. Especially since many of its features are currently being implemented, it would behoove the Department to explore similar moves for other product types. Satellites, sensors, and engines are each excellent potential opportunities for a second phase of LLP.

Section VIII: Lessons Learned: How to Spur an Acquisition Renaissance

The zeitgeist behind LLP has been framed across essays and books that carry evocative, often poetic titles. They can be categorized into historical accounts (e.g., Paul Kennedy’s *The Arsenal of Democracy* and Maury Klein’s *A Call to Arms*), diagnoses of the current situation (e.g., Karp and Zamiska’s *The Technological Republic*, Jones’ *The American Edge*, Shah and Kirchhoff’s *Unit X*), and more forward-looking books that prognosticate future policy choices (e.g., Sankar and Hart’s *Mobilize* and Freymann et al.’s *The Arsenal of Democracy*). That there are two recent works that share the ‘arsenal of democracy’ shibboleth, yet exist on different time horizons, serves as a useful coda for this essay. To put this essay’s motivating problem in simple terms, how do assemble a second arsenal of democracy from the ashes and legacy of its 20th-century forerunner?

While the cost of rearmament is substantial—in terms of absolute fiscal expenditure as well as macroeconomic crowding out—it is a bargain if the alternative is a great power war. The United States can, and must, leverage its ‘fourth arm’ of defense: the potent combination of economic strength, financial reserves, and latent industrial capacity. Effective deterrence and, if that fails, peace on our terms, requires the Department to procure, manufacture, and field critical materiel at scale. Munitions are a *sine qua non* of 21st-century war: Large Lot Procurement is an indispensable tool to achieve an American acquisition renaissance.

Bibliography

Avio USA. (2025, December 11). *Avio USA selects Virginia for SRM manufacturing facility*.

Bowden, M. (2024, December 17). *The crumbling foundation of America's military*.

Culclasure, J. R. (2013, June 3). *Book review: Freedom's forge: How American business produced victory in World War II*.

Eaglen, M., & Harrison, T. (2025, September 05). *The US military's missile gap isn't going away*.

Fletcher, Z. B. (2025, August 4). *Pentagon awards $7.8 billion in missile contracts for US and allies*.

Frazier, A. (2025, October 25). *Fleet at risk as Navy struggles to refill missile stores*.

Green, D. (2025, February 27). *The U.S. military faces a critical missile gap against China*.

Greenway, R., Fein, J., Stern, R., Beaver, W., Doan, M., Greszler, R. . . . Peters, R. (2025). *A strategy to revitalize the defense industrial base for the 21st century.*

Hadley, G., & Tirpak, J. A. (2025, August 1). *Pentagon awards $7.8 billion in contracts for hundreds of new missiles*.

Hayet, F. A. (1945, September). The use of knowledge in society. *The American Economic Review, 35*(4), 519–530.

Herman, A. (2013). *Freedom's forge: How American business produced victory in World War II.*

Jones, S. G. (2023, February 22). *The U.S. defense industrial base is not prepared for a possible conflict with China*.

Judson, J. (2024, September 17). *Norway's Kongsberg to open new Virginia missile production plant*.

Judson, J. (2025a, January 15). *Raytheon aims to boost SM-3 missile production rates*.

Judson, J. (2025b, July 17). *Top NATO commander rushing to deliver fresh Patriots to Ukraine*.

Judson, J. (2025c, September 03). *US Army awards Lockheed record $9.8 billion missile contract*.

Kiesling, L. (2015). The knowledge problem. In C. J. Coyne & P. Boettke (Eds.), *The Oxford handbook of Austrian economics* (pp. 45-64).

Kington, T. (2025, September 25). *Italy's Avio is on track to make Standard Missile motors for Raytheon*.

Kongsberg Defense and Aerospace, Inc. (2026, January 19). *Groundbreaking for Kongsberg's missile manufacturing and maintenance facility in James City County.*

Leonard, C. (2025, October 27). *America doesn't have enough weapons for a major conflict. These workers know why.*

Lockheed Martin Corporation. (2026, January 6). *Lockheed Martin and Department of War advance landmark acquisition transformation to accelerate PAC-3 MSE production*.

Losey, S. (2024, September 30). *Air Force awards Lockheed $3.2B multiyear missile contract*.

Lowell, H. (2025, July 8). *US only has 25% of all Patriot missile interceptors needed for Pentagon’s military plans*.

McCoy, C. A. (2025, November 25). *Forging America’s 21st century defense industrial base: Applying lessons from the arsenal of democracy to modern great-power competition*.

McGinn, J., & Cook, C. R. (2025, December 05). *Putting the industrial base on a wartime footing*.

Miller, E. R. (2026, February 17). *Pentagon expands THAAD interceptor production with L3Harris deal*.

O'Hanlon, M. E., & Rocha, A. (2024, June 20). *Strengthening America's defense industrial base*.

Raytheon. (2024, July 23). *RTX's Raytheon partners with AVIO to build a more resilient U.S. defense industrial base for solid rocket motor production*.

Raytheon. (2025a, September 24). *RTX's Raytheon and Avio USA expand collaboration to accelerate Mk 104 rocket motor production*.

Raytheon. (2025b, November 10). *RTX's Raytheon, Avio sign MoU to establish new solid rocket motor facility in the U.S.*

Rodriguez, R. (2010a). *Large lot procurement - Raytheon missile procurement - Raytheon visit.*

Rodriguez, R. (2010b, August). *Large lot procurement (LLP) 3 star briefing*.

Rodriguez, R. (2023). *Large lot procurement MYP concept (pilot).*

RTX. (2026, February 4). *Raytheon partners with Department of War on five landmark agreements to expand critical munition production*.

Skyrms, B. (2001). The stag hunt. *Presidential Address.*

Soliman, M. (2025, October 1). *America’s scale problem*.

U.S. Department of War. (2025, July 31). *Contracts for Jul. 31, 2025*.

U.S. Department of War. (2026, January 6). *Department of War establishes new acquisition model to more than triple PAC-3 MSE production in partnership with Lockheed Martin*.

Wasser, B., & Sheers, P. (2025). *From production lines to front lines: Revitalizing the U.S. defense industrial base for future great power conflict*.

A Strategic Investor: Federal Equity Investing Approaches to Advance National Security

Sam Moyer—is a Research Fellow at the National Defense Industrial Association's Emerging Technologies Institute (ETI), where he focuses on private investment trends, supply chain analysis, and technology policy. Prior to ETI, he was a management consultant and researcher at Gartner, a research and advisory firm. He has worked for more than a decade in defense technology policy, with a focus on innovative contracting, small business, and technology transition. He holds a BA in Economics from the College of William & Mary and an MBA from the University of Maryland. [smoyer@ndia.org]

Abstract

The U.S. government has begun purchasing equity in private companies as a tool for advancing national security. This approach can be appropriate for a narrowly scoped set of conditions: gaps in the capital markets for investments requiring large volumes of capital, with high levels of risk, for which debt is inappropriate, and the need to govern firms in sensitive industries where existing control mechanisms are insufficiently robust.

This complex and powerful form of industrial policy is unfamiliar to most national security professionals. Yet, if government equity activities continue to expand, the acquisition workforce will increasingly encounter these instruments. This paper argues that although equity investing can advance national security, the risks are also acute, and managing those risks will require a major investment in training, governance, and control mechanisms.

Introduction: The Potential Role of Government Equity Investing

There is growing interest by national security leaders and policy-makers in industrial policy tools that can expand domestic industrial capacity, control supply chain risks, and shape the global economic landscape in strategic sectors such as semiconductors, critical minerals/materials, advanced manufacturing, and defense-unique supply chains. At the same time, there is increasing recognition that America's strong, deep, and sophisticated capital markets can contribute to national security if appropriately engaged.

Government equity investing—that is, direct ownership stakes in companies by the U.S. federal government—has emerged as one potential tool to help achieve these goals. Recently executed examples of government-involved equity transactions include Intel Corporation (electronics), MP Materials (rare earth minerals and magnets), and U.S. Steel (steel production).

Empirical research suggests the existence of a gap in capital markets for high-risk, capital-intensive scale-up projects, such as "first-of-a-kind" manufacturing centers—even when expected return on investment (ROI) is high (Institute for Security and Technology, 2025). These investments, which can require hundreds of millions or even billions of dollars, are often the logical next step for growing companies after the proof-of-concept or pilot stage, after they have outgrown venture capital or the Small Business Innovation Research (SBIR) program.

One reason for this gap is that debt markets, which commonly invest in capital-intensive industries, typically make lending decisions based on credit risk (default risk), rather than expected value of a positive ROI. Underwriting by expected value has been shown in the academic economics literature to be challenging or even impossible for lenders due to adverse selection effects that emerge when lenders raise rates to compensate for risk (Stiglitz & Weiss, 1981).

The result is an investment "valley of death" for certain projects which are ROI-positive but have a high risk of default. Many strategic projects fall into this category—especially new manufacturing projects. In this setting, government equity may be more effective at providing the necessary "patient capital" than the private capital markets or government-sponsored loans, loan guarantees, or other debt-based industrial policy mechanisms.

Another advantage of equity is the wide range of governance mechanisms it can provide to shareholders, including voting rights, information rights, and covenants. These mechanisms may be more comprehensive than other control mechanisms like the Committee on Foreign Investment in the United States (CFIUS).

From a federal fiscal health perspective, equity can be balance-sheet neutral for the taxpayer, with a potential ROI when compared with mechanisms like grants and contracts. Especially when financing a major industrial scale-up effort—which may cost billions of dollars—grants and contracts could more rapidly strain the federal budget.

Defining Equity Investing

"Equity" refers to a family of investment securities that have shared features. Among these features is the ability to receive a share of future profits, or upside, from a firm. Common types of equity include preferred stock, convertible notes, warrants, and common stock. Each of these securities constitutes a legal contract between the security owner and the firm and confers a distinct bundle of rights and obligations that can affect a company's incentives and governance over time. These factors must be negotiated and then enshrined in a set of legally binding documents, typically some combination of an investment agreement, a corporate charter amendment, stockholders agreement, a contingent proxy, and employee rights and confidentiality agreements (Sherman, 2012).

Some of the commonly defined rights and obligations include the following:

- **Economic Rights:** Equity holders typically obtain some form of economic rights, commonly in the form of financial upside from company profitability (i.e., residual claimant rights). Other economic rights include dividend payments from the company, and prioritization or "seniority" in financial payouts (i.e., liquidation preferences).
- **Control Rights:** Equity holders often gain the ability to vote as part of corporate governance. This governance role typically provides either a seat on the board of directors or a vote in board seat elections. Governance rights include veto rights over certain actions defined in the corporate charter (e.g., mergers and acquisitions [M&A] activity) and the ability to amend the corporate charter. Of particular importance to national security, equity holders can obtain veto rights over M&A and intellectual property dispositions, which can help address adversarial capital concerns.
- **Information Rights:** As owners of a firm, equity holders also gain access to critical, often confidential, company information pertaining to financial and operational performance. This information is typically shared to assist shareholders with conducting corporate governance (e.g., voting on key questions).
- **Transfer Restrictions:** Equity agreements often create limitations on who can buy shares or control the company later. Terms such as anti-dilution or "ratchet" clauses can be used to restrict future actions by the company that might reduce a shareholder's rights to economic upside or other rights.
- **Enforcement Rights:** Clawbacks, step-in rights, and affirmative or negative covenants create ongoing requirements in areas such as compensation, accounting practices, leadership requirements, and many others. Of particular importance to national security, equity holders may hold companies to pre-agreed cybersecurity and other security requirements.

In the private sector, the bundle of rights corresponding to each equity transaction is negotiable. As the Pentagon participates more in equity markets, it will be required to thoughtfully negotiate the terms of each transaction to ensure the rights it obtains are optimal for the specific goals it aims to achieve in each deal.

In practice, equity transactions can be grouped into categories that share similar features:

- **Preferred Stock:** Typically have priority to common stock in payback. Preferred shareowners often have more control rights (e.g., board seats), translating to greater influence over strategic decisions such as M&A and intellectual property (IP) disposition.
- **Common Stock:** Equity with lower priority for payback compared to preferred shareholders. Common stockholders typically have privileges to vote for the board.
- **Warrants:** Securities that provide an option to purchase shares at a set price at a later date. This provides the potential for financial upside without control or information rights.
- **Convertible and Simple Agreement for Future Equity (SAFE) Structures:** Securities that provide options for future economic rights dependent on future conditions.
- **Controlling or Dual-Class Shares:** These unusual structures can be negotiated to include ongoing control and information rights without economic rights. At the extreme, governments can receive so-called "golden shares," a term for an arrangement that provides complete control of companies by offering an automatic override in any decision put to a board for a vote (Thomson Reuters, 2026).

Current Department of Defense Approach to Equity Investing

The U.S. government has recently been piloting new approaches to acquiring company equity. As of 2026, the federal government has acquired equity in at least eight firms, including Intel, Westinghouse, and xLight (Murphy et al., 2024). The Department of Defense (DoD) specifically has sponsored equity-type transactions with at least five companies (Shivakumar et al., 2026):

- MP Materials
- Korea Zinc
- ReElement Technologies
- Trilogy Metals
- L3Harris

The government has used a range of authorities and funding types to transact equity investments, including Title III dollars and the purchase authority of the Defense Production Act (DPA). Specifically, the DPA's purchase authority to expand the defense industrial base has been used to authorize purchases of shares (MP Materials, 2025).

In one recent example, the DoD acquired stock in MP Materials, a rare earth mine, refining, and magnet production company (DoD, 2025b). The MP Materials transaction illustrates how the government is combining demand guarantees, debt support, and equity-like instruments to pursue strategic industrial outcomes.

Case Study: MP Materials

In July 2025, MP Materials and the DoD entered into a multipart transaction involving equity to create a domestic rare earth mining and magnet production supply chain (MP Materials, 2025).

One component of this transaction was the provision of strong demand signals, in the form of offtake and price floor agreements, for finished and intermediate product from MP Materials. This near-guaranteed profit arrangement enabled the DoD's Office of Strategic Capital, as well as private investment banks, to syndicate over $1 billion in debt in support of this transaction.

The DoD also received $400 million in convertible preferred stock in MP Materials, as well as a warrant to purchase additional shares in the future. As disclosed in an 8K filing, the deal did not include voting rights, board seats, or other direct control mechanisms.

Based on publicly available disclosures, which may be incomplete, the transaction was negotiated to include several rights and obligations of relevance for national security. Covenants include binding requirements to construct or expand facilities involving mining, refining, chemical processing, and magnet manufacturing, at targeted activity levels. Additionally, there are numerous prohibitions against conducting business in prohibited jurisdictions, or with non-U.S. or prohibited persons. Examples of controlled activities include M&A, major asset sales, equity transactions, and magnet sales.

If these covenants are violated, the Department can demand immediate payback of loans, seek damages, or terminate its offtake and purchase agreements.

The MP Materials deal demonstrates how equity-type arrangements can support key national security goals: in this case, domestic rare earth production capabilities. The transaction blended government demand signaling tactics, commercial-style debt and equity investments, as well as tailored control mechanisms. Used in combination, these mechanisms can support one another and shift risks and financial burdens to appropriate stakeholders.

However, transactions of this complexity also introduce risks, some of which are familiar, while others are new. The sections below examine those risks and the governance frameworks needed to manage them in this new setting.

Key Risks

Government purchases of equity-type instruments present a range of risks to the government, taxpayer, and warfighter. The U.S. acquisition workforce is acclimated to managing most of these risks in the course of normal acquisitions. However, because equity-type instruments can be highly complex, powerful, and valuable, many of these risks may manifest in a new or more profound manner. Moreover, the fact that equity investing exposes the government to financial losses, and opportunities to control private company decisions, creates novel challenges.

1. **Soft-Budget Constraints:** When the government becomes a shareholder of a firm, it can be incentivized to keep funding a weak firm to avoid realizing a loss on its invested capital or to avoid national security challenges. The concept of the soft-budget constraint was examined extensively by the Hungarian economist János Kornai, who noted that in communist Hungary, "chronic loss-makers among [government supported firms] were not allowed to fail. They were always bailed out with financial subsidies or other instruments" (Kornai et al., 2003).

Kornai noted that, in many cases, when important firms are outcompeted and begin to make losses, it can lead to "social unrest and political tension." When the government owns firms, it may encounter scenarios where it is forced to choose among competing priorities, such as eliminating jobs, taking a financial loss, or promoting national security.

2. **Reduced Innovation:** When the government supports a company through equity investing, it may inadvertently suppress innovation among other companies. Through innovation, competing companies may generate alternative, or cheaper, products when compared to the subsidized firm. A subsidized incumbent, however, would be difficult to displace — particularly when soft-budget constraints might allow it to absorb losses that the private competitor could not. If innovative firms are outcompeted by subsidized firms, this would represent a stifling of private sector innovation.

3. **Crowding Out Effects:** The DoD's equity investments depend on appropriated funds furnished by the U.S. Treasury. Such funds are derived either from taxation or debt issuance. Either of these activities removes capital from the economy that could otherwise be used for private investment. This phenomenon is known as the "crowding out" of private investment activity by the government. As a result, a government equity investment may come at the cost of a potential alternative investment.

   A similar effect can occur in the labor market. The U.S. faces a shortage of skilled workers, such as engineers. If engineers are drawn into working for a firm that the DoD has invested in, those workers will be removed from other, potentially more productive opportunities. For example, if skilled engineers are hired to construct a domestic semiconductor fab, those workers become unavailable to work in other strategic industries, such as quantum computing.

4. **Workforce and Education Gaps:** Absent a skilled, expert team, the government risks paying too much, negotiating weak or inappropriate control rights, or failing to enforce conditions. In certain scenarios, the government could overpay for equity or unfairly compete with private investors.

   When the DoD conducts equity investments, it must rely on the acquisition community to assemble and execute those deals—a community which is currently not typically trained in these techniques. Engaging with the relevant financial intermediaries, such as investment banks, and negotiating with companies requires a deep understanding of the typical business practices in the financial services industry, types of investment capital, legal aspects of investing, valuation techniques, and commercial accounting practices.

5. **Politicized Decision-Making:** After deals close, government personnel will be responsible for exercising equity rights received, such as information or control rights. If the government takes voting rights or board seats, a key risk is that the government may not be a competent board member, making sound decisions about corporate governance leading to profitability. The DoD generally lacks these skills in its acquisition workforce and faces barriers with hiring or training individuals that do.

   Control rights provide the government with mechanisms to monitor and address risky behavior in companies such as adversarial acquisitions or sensitive IP disposal. The same rights can give shareholders significant influence over major corporate actions and, in some cases, indirect leverage over management's day-to-day choices; this power can easily be abused or mismanaged when the government is a shareholder. As a shareholder, the government could be in a position to harm the business prospects of a company and undermine its purpose for obtaining equity in the first place.

6. **Mandate Creep:** When governments invest directly in firms, the funds can be used for a relatively wide range of activities compared to more constricted grants and contracts. Private investment firms typically develop explicit "fund mandates" that restrict investments to specific industries, firm types, or regions. In a government context, the mandate function may be subject to "mandate creep," with changing purposes after a change in government (e.g., an election) or a perceived change in market conditions. As a result, it can be difficult to maintain accountability of government equity programs or to judge their success via predetermined metrics or other criteria.
7. **Potential Misappropriation of Public Resources:** Like all industrial interventions, such as grants, procurement, or credit enhancements, equity investments must be awarded through competitive processes and in line with procurement integrity standards. Because of the novelty and complexity of equity investing, as well as the challenges associated with setting clear success metrics and conducting oversight, these programs can be vulnerable to misappropriation.

   Globally, there are many examples of abuse involving equity investing. For example, Malaysia's 1MDB Sovereign Wealth Fund engaged in intensive equity-based industrial policy. A corruption investigation revealed that the program had outsourced significant elements of the investment activities to financial intermediaries, who were then able to misappropriate more than $3.5 billion in taxpayer funds (Parliament of Malaysia, 2016).

Recommended Best Practices for Using Equity Mechanisms

The risks above create unusual scope for abuse and therefore require careful implementation of mechanisms to control risk. Fortunately, there are numerous well-designed programs, both in the private and public sectors, that can be drawn from. The following recommendations are designed to address one or more of the risks described above. This is not intended to be a comprehensive list of best practices. National security leaders, acquisition professionals, and policy-makers should consider opportunities to implement these tactics to control key risks.

1. **Workforce and Institutional Capability:** To successfully run an equity purchasing program, the government will need to staff its execution offices similarly to commercial investment offices. Key capacity requirements include:
   a. **Deal underwriting:** conducting financial analysis (e.g., valuations, weighted average cost of capital) and negotiating securities transactions
   b. **Portfolio monitoring:** reviewing company information, voting in board meetings, enforcing covenants
   c. **Risk management:** performance tracking, concentration limits, high-risk negotiating (e.g., disposing of sensitive or strategic assets, moving production overseas)

   To execute on a similar mission as part of the CHIPS program, the Department of Commerce recruited experienced investment professionals from companies like Goldman Sachs, KKR, and Blackstone (Vorland, 2025).

   The Department may find it necessary to structure such a cell within a Federally Funded Research and Development Center or University Affiliated Research Center in order to provide the hiring and compensation structures necessary to access the relevant workforce. Of particular importance will be flexibility on work location, since most investment professionals are concentrated in large investment centers like New York City, instead of Washington, D.C. or other centers of government activity. Work-from-home options, or physical offices in those financial centers will be essential to recruiting skilled personnel and for the mission.

2. **Minimally Intrusive Control Rights:** Controlling stakes, golden shares, or board seats should be exceptional tools used only with explicit statutory grounding and clear necessity. In many cases, security concerns about foreign influence are better addressed by non-equity tools such as

    a. CFIUS enforcement and conditions

    b. Export controls

    c. Government contract supply chain requirements

    When use of control mechanisms is deemed warranted, the DoD should be prepared to negotiate for minimally invasive control rights and remedies. Control rights must be implemented as a mechanism to either prevent or raise awareness when a narrow set of highly consequential activities take place, such as moving IP or production offshore. Rather than provide a blanket veto over company decisions, the DoD equity should instead aim to offer the government some degree of voice and visibility with respect to decisions that could negatively impact national security.

    In general, the types of control rights and remedies that are commonly provided to private preferred equity holders may offer an appropriate template. Common control provisions for preferred equity holders include veto rights over M&A and IP dispositions, as well as covenants requiring high standards for security.

    The 2025 MP Materials transaction illustrates this approach: the government received non-voting preferred equity and a warrant, with covenants (rather than voting rights) controlling for key risks such as foreign ownership, customer selection, and asset disposition.

    These control mechanisms can also be further limited to specific time periods or geographic regions. Board observer seats, or mutually selected independent board members are also common mechanisms for shareholders to exercise influence in a balanced manner.

    Since these terms are common in industry, they may also be perceived as less intrusive than other formulations—such as arrangements like golden shares.

3. **Oversight Mechanisms:** Equity investing has generally been carried out under the aegis of the DoD's industrial policy toolkit, such as Title III Fund (Harrell, 2025). The DoD currently collects information about its industrial policy activities via the National Defense Industrial Strategy, which reports on overall progress across several qualitative levels of effort for strengthening the industrial base, such as "onshoring critical production capacity." The DoD also reports in ad hoc briefings on its use of the DPA. The Office of Strategic Capital is required to submit an annual report.

    Equity investments are powerful industrial policy actions with the potential to have broad effects on the economy and society. Proper oversight is required to provide accountability to Congress and transparency to the taxpayer. At the same time, Congress is not well-positioned to oversee day-to-day investment decision-making. Considerations that must be balanced include data privacy (including specific deal terms, as well as information provided to the government as part of company board reporting), ethical oversight, investment performance metrics, and other non-pecuniary performance metrics. The oversight mechanism should be designed to allow Congress to assess the financial performance and strategic impact of the DoD's equity investments in a timely manner and to mandate timely reporting of urgent updates involving risk. Important information to be transmitted should include the following:

a. **Financial Statements:** A balance sheet and other financial statements reflecting the financial performance of the government's equity positions. Data should include the current fair value of the DoD's equity portfolio and the amount paid for those assets.

b. **Accounting/Scoring:** An explanation of how and when equity stakes are valued.

c. **Fund Performance Metrics:** Relevant information on the activities and results of government equity investing activities (National Defense Authorization Act, 2025b). This could include default and recovery rates, private capital mobilized, and ROI.

d. **Risk Management Considerations:** Portfolio concentration and performance against anticipated results.

e. **Impact on DIB**: Impacts on small businesses (e.g., financial fragility metrics, supply chain resiliency impact), technology performance milestones, NAICS code, or investment locations.

This information could be provided to Congress across a range of form factors, such as a public annual report, a classified annual report, or in-person briefings. Congress should consider expanding these to include an "investment scorecard" for information on how equity investments are being managed across the DoD.

Ad hoc briefings may also be required if triggered by certain events (e.g., investments above some threshold size). For instance, the U.S. Development Finance Corporation (DFC) provides prenotification to the congressional foreign affairs committees if a deal over $20 million is intended to be transacted (22 U.S.C. § 9656). Ad hoc briefings may also be helpful in scenarios involving the acquisition of majority ownership of firms, controlling stakes, or major breaches of covenants.

4. **Internal Governance:** Given the compliance and ethics risks outlined above, programs engaging in government equity investing will require thoughtful internal governance systems. Strong governance can protect the taxpayer from abuse and ensure satisfactory transparency about significant industrial policy decisions. Strong internal guardrails will be required to prevent misuse of authority and ensure accountability.

   Among private investment firms, a critical governance role for equity investing is often the investment committee (IC). The purpose of this committee, often composed of senior investment professionals, is to ensure each investment is conducted in accordance with an agreed investment policy, such as the investment mandate (Miranda & Corcoran, 2026). For governance purposes, the committee provides a forum for making sound investment decisions, as well as a written record of how investment decisions were made.

   In the private sector, a chief compliance officer (CCO) is typically considered a requirement for organizations conducting complex financial transactions. The Department of Justice's Evaluation of Corporate Compliance Programs (ECCP)—the authoritative guidance federal prosecutors use to assess compliance adequacy—calls for a designated CCO with genuine independence (defined partly as maintaining a separate line of reporting) from the deal teams they oversee (Department of Justice, 2024). A DoD equity program should designate a CCO, ombudsperson, or comparable official with a direct reporting line to senior leadership and sufficient resources for ongoing monitoring. In addition, particularly for complex financial issues, a chief audit executive, which can provide a strong and thorough internal audit function and facilitate independent outside audits, is standard.

Another common approach is inspector general offices, independent oversight officials established across the federal government with authority to conduct audits, investigations, and evaluations of their host agency's programs and operations. The DFC created a dedicated Inspector General that conducts both audits (covering financial statements, cybersecurity, and payment integrity) and discretionary performance audits and reports regularly to Congress, the DFC Board of Directors, and DFC management (International Development Finance Corporation, 2026).

5. **Portfolio Risk Management and Exit Policy:** Among private investment firms, a range of sophisticated risk management tactics are used to prevent negative portfolio outcomes. Leading asset managers employ dedicated quantitative teams, proprietary models, and real-time monitoring systems to manage risks such as outsized losses, illiquidity, and overconcentration. Selected methodologies include

    a. **Stress testing and scenario analysis:** Asset management firms often use stress testing to evaluate how well portfolios would perform under extreme but plausible scenarios, including financial crises or changing interest rates (Kenton, 2023). In the private sector, such stress testing is often required by law.

    b. **Liquidity risk management:** Investment firms often use models of portfolio liquidity, such as "time to liquidation" or "market depth" to design portfolios (BlackRock Investment Institute, 2014). This may be relevant for government equity investments, which could concentrate in capital-intensive sectors. If an asset cannot be sold without significant value loss, exiting that asset may be painful for the owner, making continued support more likely.

    c. **Concentration limits and exposure caps:** Investment firms often maintain "exposure limits" to prevent excessive investment in single companies, sectors, or regions (AssetVantage, 2026). Portfolio rebalancing may be required when exposure limits are reached.

In practice, the application of portfolio risk management mechanisms like these can involve complex mathematical analysis. Applying these methodologies in practice can require expertise as well as experience managing large portfolios of assets in related sectors. Managing the risk of an equity portfolio in accordance with industry best practices will require the development of a tailored risk-management approach by an expert team.

Conclusion

U.S. national security increasingly depends on economic and industrial resiliency. Where traditional acquisition tools struggle—supporting capital-intensive projects in sensitive industries—equity instruments may have a unique role.

However, equity investments also create novel opportunities for abuse by government due to their complexity, power, and challenging incentives. To take advantage of this tool, the DoD will need to invest in its own acquisition and oversight capabilities and implement proven best practices to control risks.

References

AssetVantage. (2026, January 23). *Look-through analysis: Tracking concentration risk with mutual fund investments.* https://www.assetvantage.com/blogs/look-through-analysis-tracking-concentration-risk-with-mutual-fund-investments/

BlackRock Investment Institute. (2014, July 3). *The liquidity challenge.* https://www.blackrock.com/sg/en/market-insights/article/sg-one/DEFAULT/investment-insight/the-liquidity-challenge

Department of Justice. (2024, September). *Evaluation of corporate compliance programs.* https://www.justice.gov/criminal/criminal-fraud/page/file/937501/dl

DoD. (2025a, April 1). *Office of strategic capital receives $8.9 billion in financing requests for first domestic manufacturing loan program.* https://www.war.gov/News/Releases/Release/Article/4141755/office-of-strategic-capital-receives-89-billion-in-financing-requests-for-first/

DoD. (2025b, August 10). *Office of strategic capital announces first loan through DoD agreement with MP Materials to secure critical materials supply chain.* https://www.war.gov/News/Releases/Release/Article/4270722/office-of-strategic-capital-announces-first-loan-through-dod-agreement-with-mp/

Harrell, P. E. (2025, August 28). *The legal bases for government stakes in private firms.* Lawfare. https://www.lawfaremedia.org/article/the-legal-bases-for-government-stakes-in-private-firms

Institute for Security and Technology. (2025, October). *The missing middle: How to close America's deep-tech financing gap in strategic competition with China.* https://securityandtechnology.org/virtual-library/report/the-missing-middle/

International Development Finance Corporation, Office of Inspector General. (2026, March 25). *Office of inspector general.* https://www.dfc.gov/oig

Kenton, W. (2023, June 29). *Stress testing: Techniques, purpose, and real-world examples.* Investopedia. https://www.investopedia.com/terms/s/stresstesting.asp

Kornai, J., Maskin, E., & Roland, G. (2003). Understanding the soft budget constraint. *Journal of Economic Literature, 41*(4), 1095–1136.

Miranda, S., & Corcoran, B. (2026). *Investment committee best practice.* Partners Capital. https://partners-cap.com/insights/investment-committee-best-practice/

MP Materials Corp. (2025, July 10). *Current report on Form 8-K*. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1801368/000119312525157310/d43796d8k.htm

Murphy, E. L., Cook, C. R., Harding, E., et al. (2024, November 21). *Alternative funding mechanisms in review.* Center for Strategic and International Studies. https://www.csis.org/analysis/alternative-funding-mechanisms-review

National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119–60, § 867 (2025a).

National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119–60, § 8755 (2025b).

Parliament of Malaysia, Public Accounts Committee. (2016). *Report of the public accounts committee on 1MDB.*

Sherman, A. J. (2012). *Raising capital.* Amacom Books.

Shivakumar, S., Wessner, C., & Tutino, C. (2026, February 12). *Understanding federal equity investments in strategic companies.* Center for Strategic and International Studies. https://www.csis.org/analysis/understanding-federal-equity-investments-strategic-companies

Sidley Austin. (2022, June 23). *U.S. DOJ's compliance certifications put chief compliance officers in criminal crosshairs.* https://www.sidley.com/en/insights/newsupdates/2022/06/us-dojs-compliance-certification-policy-puts-chief-compliance-officers-in-criminal-crosshairs

Stiglitz, J. E., & Weiss, A. (1981, June). Credit rationing in markets with imperfect information. *American Economic Review, 71*(3), 393–410.

Thomson Reuters. (2026, March 20). *Golden share.* Practical Law. https://uk.practicallaw.thomsonreuters.com/4-107-6657

U.S.Code § 9621 (BUILD Act).

U.S. Code § 9656.

Vorland, D. (2025, December 3). *How private sector experience helped the CHIPS Act succeed.* Blue Horizon Project. https://pub.bluehorizonproject.org/p/how-private-sector-experience-helped

Strengthening the Arsenal of Deterrence: Modernizing the DoD Munitions Requirements Process to Include Allied Equipping

Sara Eighmey—is a Department of Defense professional with experience supporting joint and defense acquisition processes related to operational requirements and strategic planning. Her work focuses on the intersection of deterrence, munitions planning, and industrial-based considerations, including demand forecasting, risk management, and resourcing trade-offs. She is a graduate student whose research examines how U.S. defense processes can better align with real-world warfighting demands, emerging policy guidance, and evolving security dynamics in the Indo-Pacific region. [sara.j.eighmey.civ@us.navy.mil; sara.eighmey@nps.edu]

Abstract

The Department of Defense Munitions Requirements Process provides the analytical foundation for identifying and prioritizing joint force munitions needed to execute approved war plans, yet its historical focus on U.S.-only operational demand limits its ability to account for the increasingly coalition-dependent character of contemporary conflict. This paper argues that the existing framework is insufficient and should be expanded to incorporate two critical dimensions: allied and partner equipping, and co-production of key munitions with trusted foreign defense industries.

Drawing on a comparative assessment of established analytical methodologies and emerging policy guidance on industrial base resilience and partner capability development, this paper identifies seams and statutory constraints that complicate integrating allied and partner considerations into munitions planning. It evaluates how incorporating allied and partner demand signals and co-production capacity could affect demand forecasting, risk/burden sharing, and resourcing trade-offs within resource allocation cycles.

The analysis also examines the distinct strategic considerations associated with defense articles provided under the Taiwan Relations Act of 1979 and emerging bilateral and trilateral initiatives, given their role in deterrence, industrial surge capacity, and regional stability. This paper concludes by proposing an expanded, coalition-integrated munitions requirements framework to improve analytical coherence, strengthen industrial resilience, and support credible collective deterrence.

Introduction: Institutional Design and the Credibility of Deterrence

Modeling and simulation outputs and recent wargames consistently reveal a recurring operational vulnerability: U.S. and allied munitions stockpiles are depleted far faster than current planning assumptions anticipate, exposing a growing gap between operational demand and industrial capacity (Center for Strategic and International Studies [CSIS], 2025). This divergence between projected requirements and available inventories raises fundamental questions about the adequacy of existing munitions planning frameworks.

Recent analyses have highlighted declining stockpile depth, limited industrial surge capacity, and extended production timelines for high-demand precision munitions (Clark et al., 2025; Fusco et al., 2025; Greenwalt & Patt, 2025). While much of the current literature emphasizes resourcing and industrial capacity constraints, far less attention has been paid to how the architecture of the requirements generation process itself shapes demand signals and conditions industrial response. In practice, what appears downstream as industrial shortfalls often originates upstream in the institutional design of planning methodologies.

This dynamic is increasingly reflected in analysis of the defense industrial base, which argues that U.S. munitions strategy has overemphasized *depth*—the ability to produce large quantities of narrow sets of systems—at the expense of *breadth*, or the diversity of munitions, suppliers, and production pathways required for resilience (McGinn, 2026). Rather than mitigating risks, a depth-centric approach can reinforce structural fragility by concentrating production and limiting the system's ability to respond to disruption or shifting operational demands.

These limitations are particularly consequential in coalition-dependent and logistically contested environments. Emerging initiatives such as AUKUS and the trilateral security partnership among Australia, the United Kingdom, and the United States illustrate how expanding production across allied networks can enhance resilience, adaptability, and operational reach (Lehn, 2026). Recent conflicts reinforce these concerns: Russia's invasion of Ukraine demonstrated how quickly precision-guided munitions and critical enablers can be depleted and exposed the fragility of global production pipelines (GAO, 2024). In the Indo-Pacific, U.S. strategy anticipates sustained, multi-domain operations against a peer adversary; a depth-centric industrial model risks proving insufficient in both scale and flexibility (Paparo, 2025).

ADM Paparo's characterization of the region's time–distance challenge further underscores the urgency of these issues, highlighting how the ability to act inside the adversary's decision cycle—through accelerated sensing, movement, and weapons delivery—directly intersects with the industrial and coalition-based reforms this paper evaluates (Paparo, 2026). As he emphasizes,

> Deterrence doesn't start at the shoreline—it starts with the ability to integrate, train, secure, sustain, and scale what the fight requires—at speed, at scale, and with allies and partners. Faster than the problem changes and faster than the competitor can adjust. And that's why we're operationalizing readiness.

This suggests that improving industrial resilience is not solely a matter of increasing production, but of redesigning the institutional mechanisms that generate and structure demand in the first place.

To understand why these vulnerabilities persist and how they intersect with coalition and industrial reforms, it is necessary to examine the historical and institutional context of U.S. munitions planning, including the intricacies of the munitions requirements-generation process and industrial base structures.

This paper asks, *How can the Department of Defense (DoD) modernize the Munitions Requirements Process (MRP) to incorporate coalition demand and allied industrial capacity while preserving the institutional separation between Title 10 operational planning and Title 22 security cooperation authorities?* Addressing this question requires examining how institutional design shapes demand signals, industrial planning, and coalition readiness in an era of coalition-dependent deterrence.

Research Design

This paper offers a qualitative policy analysis drawing on existing literature on defense industrial base resilience, munitions planning processes, and security cooperation authorities. It compares the structure of the DoD MRP with statutory frameworks governing partner equipping to identify structural seams that affect demand forecasting and industrial planning. The analysis synthesizes policy guidance, congressional reports, and defense industrial base studies to develop recommendations for a coalition-integrated planning framework.

Foundations and Enduring Limitations of the Munitions Requirements Process

The DoD MRP is the primary analytic mechanism used to estimate the quantity of munitions required to execute approved operational plans. The U.S. MRP, as it exists today, traces its formal establishment in the late 1990s (Kress, n.d.). Codified in DoD Instruction 3000.04, the MRP cycle begins with the DIA threat assessment. It is structured around the Combatant Command (CCMD) Phased Threat Distribution (PTD) conference, during which threats are allocated among the theater components. Following the PTD conference, each military service conducts independent modeling and simulation (M&S) and convenes an internal Munitions Requirements Process Review Board (MRPRB) to validate and adjust to account for planning constraints, including projected force posture, weapon effectiveness, range, and magazine depth. The service then produces a sufficiency assessment report, intended to inform its Program Objective Memorandum (POM) investments, and provides a copy to the CCMDs. The CCMD assesses the service's quantitative and qualitative data to inform their Operational Risk Assessment (ORA), which is subsequently submitted to the Joint Staff to inform the POM cycle (DoD, 2018).

Although designed to support POM decisions and enhance collaboration among CCMDs, it fails to achieve either objective effectively (GAO, 2002). The mechanics of the process are clear; the strategic assumptions embedded within it are less visible but far more consequential. The consequences of these embedded assumptions extend beyond planning mechanics—they shape how risk is defined, absorbed, and transmitted through the broader defense system.

The MRP evolved within a planning paradigm that emphasized U.S.-centric operations, predictable timelines, and available surge capacity. It is implicitly premised on linear escalation dynamics and relatively static weapons employment concepts—assumptions increasingly misaligned with contemporary warfare, which is often characterized by distributed operations, contested logistics, and rapid consumption rates (CSIS, 2024). These limitations are particularly consequential in coalition-dependent theaters, where U.S. munitions planning is closely intertwined with allied demand, interoperability requirements, and shared production capacity. They are further complicated by statutory and policy frameworks governing the provision of defense articles to partners, with the most notable in the Pacific being the Taiwan Relations Act (TRA), which generates additional operational and strategic demand signals not fully captured within existing requirements models.

Emerging technologies further destabilize forecast-based planning. Rapid adaptation of long-range precision weapons, integrated countermeasures, adaptive targeting systems, and the use of attritable systems compress technology cycles and accelerate expenditure rates. Cost asymmetries between relatively inexpensive offensive systems and more costly defensive responses compound these pressures (Black, 2025).

Greenwalt and Patt (2025) argue that requirements systems remain oriented toward a forecast-centric model of conflict that assumes stable threats and predictable trajectories. In munitions planning, this orientation incentivizes organic optimization against static scenarios while discounting rapid technology adaptation, learning competition, and accelerated expenditure rates. Adversaries operating on technology cycles measured in months rather than decades can thus explore assumptions embedded in requirements models, contributing to persistent underestimation of demand and erosion in deterrence credibility (Greenwalt & Patt, 2025).

Without structural adjustments, current munitions requirements methodologies risk underestimating both the scale and duration of consumption in high-intensity conflict. The result is not simply an inventory imbalance, but diminished deterrence credibility and reduced capacity to sustain decisive operations.

The Evolution of the U.S. Defense Industrial Base: From Consolidation to Constraint

The U.S. defense industrial base as it exists today is the product of generations, shaped incrementally as the strategic landscape evolved. It reflects decades of institutionalized demand assumptions embedded within the MRP, which in turn informed procurement and shaped production decisions. Major changes in the international security landscape have historically driven cyclical fluctuations in U.S. defense spending, marked by periods of boom and bust. While it inadequately forecasted recent shifts, it also lacked the agility to respond to the demands of modern strategic competition (Sullivan, 2024).

Figure 1 demonstrates the shrinking pool of U.S. defense prime contractors between 1990 and 2025, precipitated by the collapse of the Soviet Union, which initiated a sustained decline in companies operating across all major defense systems (e.g., fixed-wing aircraft, missiles, ships, munitions). As a result, major defense firms merged over the last two decades, falling from 51 major companies to only six large prime contractors – Boeing, Lockheed Martin, Northrop Grumman, RTX (formerly Raytheon), Huntington Ingalls Industries, and General Dynamics (CSIS, 2025).

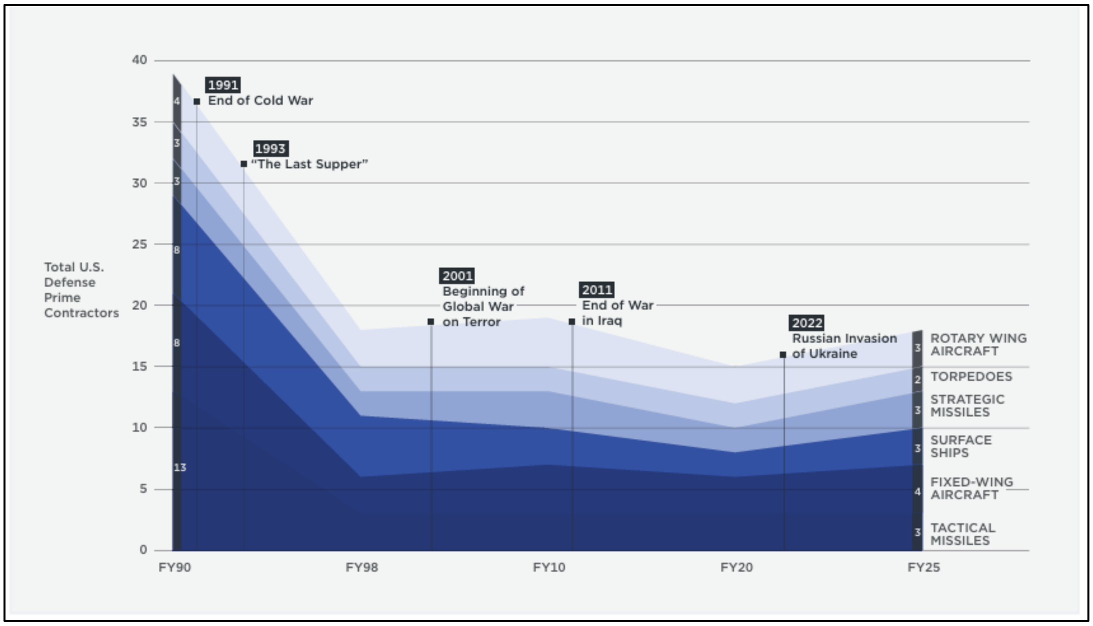

Figure 1. U.S. Defense Prime Contractors, 1990–2025 (Wasser & Sheers, 2025)

Table 1 illustrates the stark mismatch between annual and peak production rates for key U.S. munitions, revealing the extraordinary time and investment needed to procure 10,000 units even under maximum surge conditions. High-end weapons such as the SM-6 and JASSM-ER demand particularly lengthy timelines—with approximately $33 billion and 80 years for 10,000 SM-6s, and $10 billion and 19 years for the same quantity of JASSM-ERs—especially when contrasted with simpler, short-range munitions like SDBs and JDAMs.

Table 1. Estimated Production Timelines and Costs for Selected U.S. Munitions (Mitchell Institute for Aerospace Studies, 2023)

Weapon	Average annual production rate	Highest one-year production rate	Cost to acquire 10,000 weapons	Time to acquire 10,000 weapons at the highest one-year production rate
SM-6	115	125	$33 billion	80 years
JASSM-ER	257	525	$10 billion	19 years
Notional new stand-in weapon	2,500	5,000	$3 billion	2–4 years
SDB II	1,716	2,910	$1.86 billion	3 years
SDB I	2,500	6,878	$.040 billion	1.5 years
JDAM	16,780	43,594	$0.30 billion	0.2 years

Figure 2 underscores the broader strategic implications of these constraints. If "production is deterrence," then enabling Taiwan to withstand potential aggression from the People's Republic of China (PRC) requires an industrial base capable of surging munitions output not only for U.S. forces but also for allies and partners who rely on these systems to maintain deterrence postures and promote regional stability.

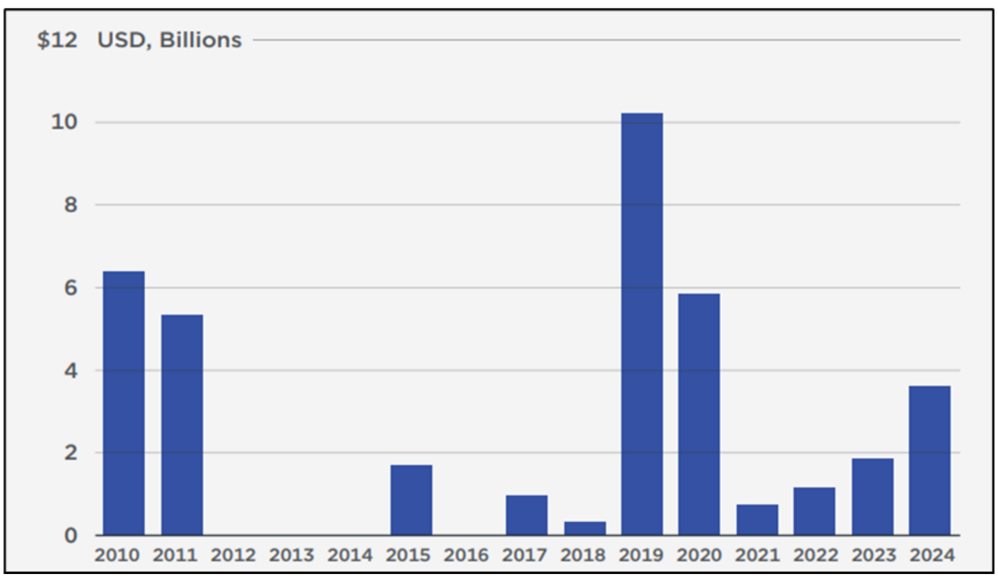

Note. Pilot training, contractor support, and select logistics services are omitted.

Figure 2. Sales to Taiwan Containing Platforms or Products Sourced from the U.S. Defense Industrial Base (Wasser & Sheers, 2025)

The RAND Commission on the National Defense Strategy (2024) notes that allies and partners are foundational to U.S. security, and effective deterrence requires integrated capabilities, interoperability, and collective resolve. The Commission highlights co-production, coordinated industrial base investments, and streamlined technology sharing to turn coalition plans into actionable capability. Achieving this depends on modernization in the Federal

Acquisition Regulation (FAR) to reduce regulatory friction, rapidly adapt commercial practices, and create predictable demand signals. This recently implemented reform—coupled with procurement practices like steady demand signals and multiyear procurement plans—will enable small and middle-tier suppliers to scale production and better support joint requirements (CSIS, 2026).

The resilience of the defense industrial base depends not only on funding levels but on the stability and credibility of demand signals generated by the requirements process. Small businesses—critical suppliers of specialized components, advanced manufacturing capabilities, and emerging technologies—rely on predictable procurement pathways and multiyear planning horizons to scale production (CSIS, 2026). Fluctuating requirements or fragmented decisions disproportionately affect these firms, eroding surge capacity and weakening the innovation ecosystem that underpins long-term military advantage.

Despite significant changes in the global security environment over the past three decades, the MRP has not fully adapted. It does not systematically account for munitions needed to support allies and key partners, nor does it incorporate co-production initiatives that could expand industrial capacity and reduce risk. This narrow focus risks miscalculating aggregate demand, undermining coalition readiness, and constraining deterrence options.

Title 10 and Title 22: Institutional Seams or Structural Fault Lines in Coalition Munitions Planning?

Under Title 10, the DoD is responsible for organizing, training, and equipping U.S. forces to execute approved war plans, a mandate codified in the statutory missions of the military departments and rooted in the post–Goldwater-Nichols division of responsibilities (United States Congress, 1986). The MRP is embedded within this construct, deriving demand from force-planning scenarios, threat assessments and physics-based modeling, producing analytically coherent, sufficiency-oriented requirements for U.S. forces (DoD, 2018). Through the Planning, Programming, Budgeting, and Execution (PPBE) process, these requirements signal the defense industrial base to sustain the operational demand generated endogenously within the DoD's force-development system.

In contrast, Title 22 authorities govern the provision of defense articles and security assistance to foreign partners. These authorities authorize foreign military sales (FMS), security assistance, and defense article transfers, including those executed under Presidential Drawdown Authority (PDA). Managed primarily by the Department of State and implemented through statutes such as the Arms Export Control Act, Foreign Assistance Act, and TRA, these mechanisms are transaction-based, politically mediated, and subject to congressional notification and diplomatic prioritization (Congressional Research Service, 2020; Kane, 2016). As a result, allied and partner demand enters the system exogenously through security cooperation channels rather than through the forecast-based modeling processes that inform the DoD's internal requirements.

Although both Title 10 and Title 22 systems are designed to mitigate risk, they operate in fundamentally different domains—operational versus political—creating structural misalignment in how demand is generated. Title 10 focuses on reducing risk in executing contingency operations, while Title 22 is designed to manage diplomatic, escalation, and political risks associated with security assistance. Despite these divergent logics, the industrial capacity that supports both systems is unitary: the same industrial base and production lines must supply U.S. operational requirements and partner-equipping programs. Although Title 10 and Title 22 authorities operate under distinct legal and bureaucratic frameworks, they ultimately rely on the same industrial base and production lines. This creates a structural condition in which operational demand and partner equipping are governed separately but remain industrially interdependent. Table 2 illustrates these institutional distinctions between Title 10 and Title 22 authorities that generate these parallel signals, while Table 1 formalizes their underlying structural framework. The resulting bifurcation manifests as bureaucratic friction but also generates parallel signals that distort industrial planning, obscure coalition requirements, and shift critical trade-offs from analytical planning processes to political decision points.

Table 2. Title 10 and Title 22 Institutional Seams in Munitions Demand Generation

Dimension	Title 10	Title 22
Purpose	Equip the U.S. Joint Force to achieve operational objectives	Enable partner capacity and support strategic deterrence
Strategic Objective	Mitigate operational risks in executing contingency operations	Mitigate diplomatic and escalation risks
Time Horizon	Forecast-based expenditures in executing approved war plans	Case-by-case transfers driven by real-time contingencies
Demand Source	Combatant Command Operational Plan (OPLAN) requirements	Partner requests and political prioritization
Budgeting	PPBE / POM processes	Security assistance and assistance accounts
Decision Authority	Military Departments and the Secretary of Defense	State Department with congressional oversight

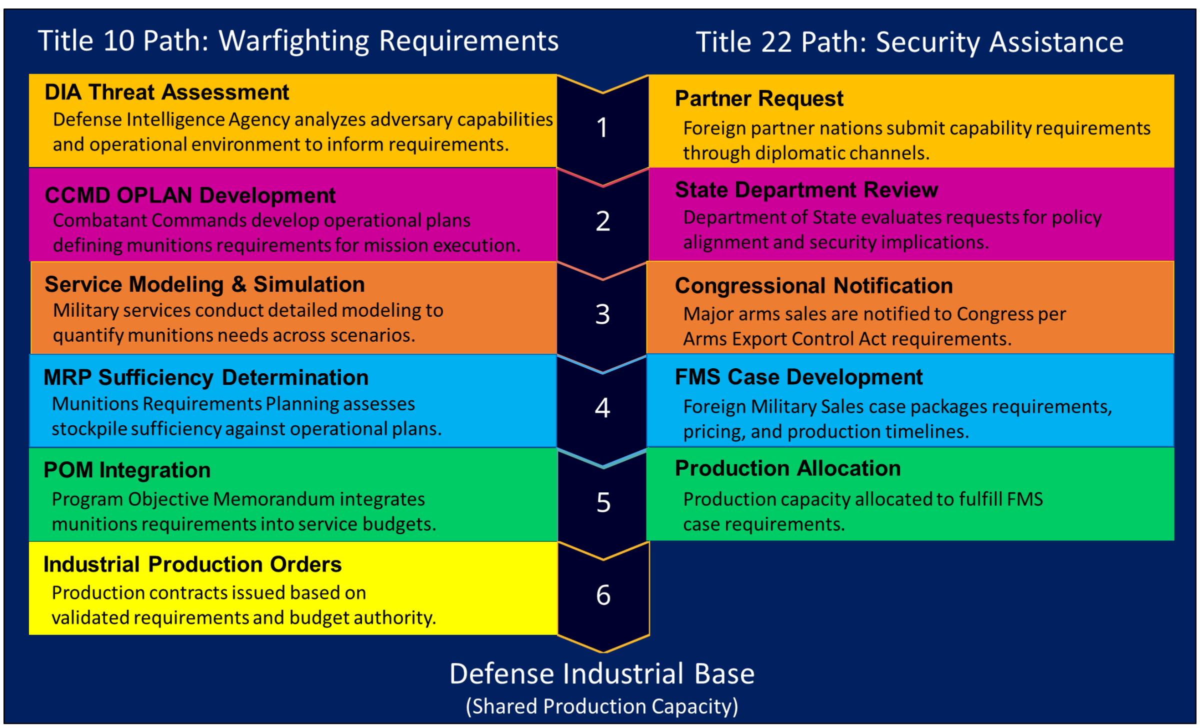

Figure 3. Title 10 and Title 22 Institutional Seams in Munitions Demand Generation

U.S. policy governing the provisions of defense articles to Taiwan illustrates the implications of this divide. The TRA directs that the United States "make available to Taiwan such defense articles and services in such quantity as may be necessary to enable Taiwan to maintain a sufficient self-defense capability" (Taiwan Relations Act, 1979). This statutory language establishes a strategic obligation rather than a discretionary or episodic policy choice. Subsequent legislation has reinforced this requirement by directing the DoD to support the acceleration of Taiwan-related defense articles, many of which are munitions-intensive and constitute a predictable and recurring demand signal.

Despite this statutory clarity, Taiwan-related munitions demand is generally treated as external to the formal MRP and is often addressed through ad hoc security cooperation mechanisms rather than integrated into demand forecasting. The absence of these requirements in core munitions planning risks understating total demand and misaligning industrial capacity with legally mandated policy objectives, reinforcing the structural disconnect between operational planning and partner equipping that characterizes the Title 10–Title 22 divide.

Rearchitecting the Munitions Requirements Process for Coalition-integrated Deterrence

The limitations of the current MRP are not solely a function of resource constraints, but of how demand itself is defined and generated. To address this structural shortfall, this paper introduces the Coalition Munitions Integration Framework (CMIF)—a demand-generation and industrial-alignment model designed to integrate U.S., allied, and partner munitions requirements into a unified analytical structure.

The CMIF departs from the traditional U.S.-centric sufficiency model by explicitly incorporating coalition demand signals and distributed production capacity into early-stage planning. Rather than treating allied equipping and co-production as external or reactive considerations, the framework embeds them within the core analytical processes that shape requirements, resourcing, and industrial planning decisions.

Conceptually, the CMIF is organized across three interdependent layers. The U.S. Requirements Core retains primacy, anchoring the framework in validated operational plans and Title 10 responsibilities. Surrounding this core, the Allied Equipping Layer integrates partner demand signals into modeling and simulation outputs, enabling planners to estimate aggregate coalition consumption requirements. The outer Co-Production Capacity Ring captures the contributions of allied and partner industrial bases, enabling a more accurate assessment of surge capacity, supply chain resilience, and distributed manufacturing potential. Together, these elements form a coalition-oriented demand and production framework (see Figure 4).

Figure 4. Coalition Munitions Integration Framework

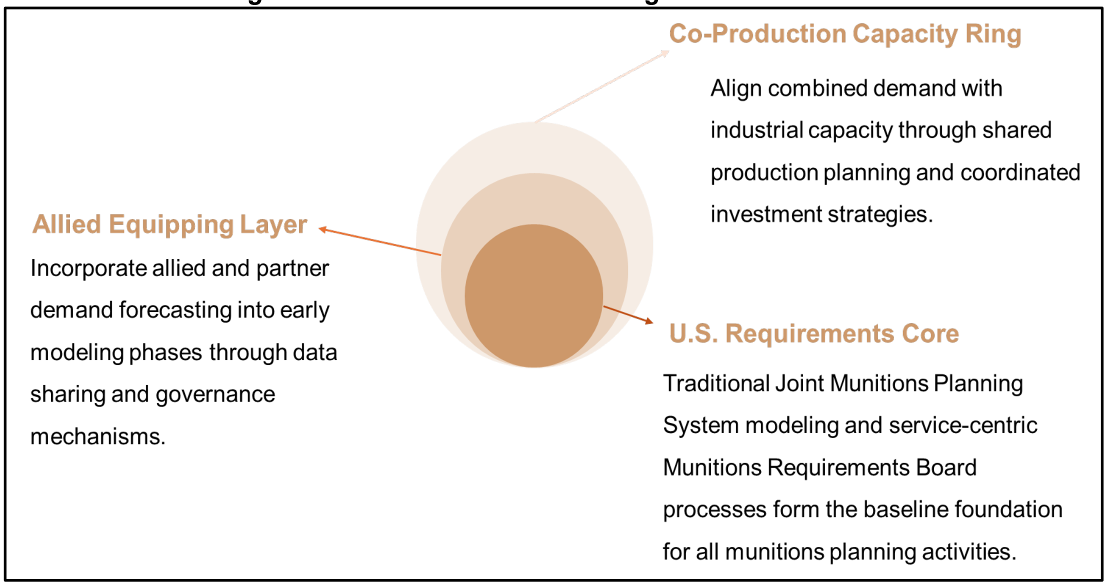

Operationalizing the CMIF requires three institutional adjustments. First, allied and partner demand signals must be incorporated into early-stage modeling inputs within the MRP cycle, enabling estimation of total coalition requirements rather than U.S.-only demand. Second, co-production arrangements with trusted partners must be integrated into industrial capacity models to reflect the realities of distributed manufacturing and allied surge capacity. Third, governance mechanisms—potentially coordinated through existing security cooperation structures—must enable structured data sharing among the services, combatant commands, and key partners.

Importantly, the CMIF does not collapse the statutory distinction between Title 10 and Title 22 authorities. Instead, it preserves their legal separation while analytically integrating their effects at the point of demand generation. In doing so, the framework reduces the current reliance on ad hoc political decision-making to reconcile competing demands on a shared industrial base.

By aligning operational planning, coalition requirements, and industrial capacity within a single analytical framework, the CMIF enables more accurate demand forecasting, more efficient resource allocation, and more credible deterrence outcomes. In contrast to the existing model—where industrial shortfalls are often discovered only after requirements are validated—the CMIF allows planners to identify and mitigate risk upstream, before crisis conditions force reactive adjustments.

Implications and Expected Outcomes

The CMIF has significant implications for accelerating credible deterrence in the Indo-Pacific and beyond while advancing U.S. strategic primacy. Current processes estimate only U.S. requirements, overlooking coalition demand and shared operational needs, creating a coalition demand gap. This gap is magnified in the Indo-Pacific by contested logistics, extended lead times, the tyranny of distance, and limited replenishment options. Expanding U.S. domestic industrial capacity to a collaborative production network with allies and partners addresses these constraints, strengthens partnerships, enhances supply chain resilience, and mitigates production bottlenecks that contribute to long-lead times and costly sustainment.

Third, the framework enables the United States to demonstrate rapid response, delivering prompt defense articles that signal commitment and strengthen deterrence. Lastly, enabling coordinated, data-driven planning enables the model to generate predictable demand signals for industry and allied partners, fostering preparedness and aligning coalition readiness with broader, shared strategic objectives. Collectively, this model institutionalizes an integrated planning framework that strengthens U.S. leadership and maximizes coalition effectiveness.

Conclusion

This paper argues that the institutional separation between Title 10 warfighting planning and Title 22 security assistance authorities produces structural distortions in munitions demand forecasting. In an era where coalition capability is integral to operational success, this separation is no longer a minor bureaucratic inefficiency, but a strategically consequential limitation.

The contemporary security environment—shaped by the People's Republic of China's rapid military modernization, emerging patterns of cooperation among revisionist states, and the proliferation and militarization of low-cost technologies—demands a more integrated approach to deterrence. As the strategic landscape evolves, the United States must leverage the defense industrial base not only as a procurement system but also as an instrument of statecraft. This requires targeted investments, predictable demand signals, and deeper cooperation with trusted allies and partners to sustain the balance of power.

In the Indo-Pacific, where deterrence depends on distributed operations, contested logistics, and coalition integration, munitions planning can no longer remain a national sufficiency exercise. Treating allied and partner requirements as external to formal planning processes risks underestimating total demand, misaligning industrial capacity, and weakening deterrence credibility.

The CMIF provides a path forward by integrating allied demand forecasting and co-production capacity into the existing MRP. By aligning operational planning, coalition requirements, and industrial realities, this framework would allow planners to model total coalition demand, manage industrial risk more efficiently, and synchronize production capacity with operational plans before crisis conditions force reactive political decisions.

Ultimately, strengthening deterrence in the Indo-Pacific requires institutional alignment between operational planning, industrial capacity, and alliance commitments. Modernizing the MRP to incorporate coalition demand signals is therefore not merely a procedural reform but a strategic adaptation to the realities of contemporary coalition warfare.

References

Arabia, C. L. (2024, November 15). *Defense primer: DoD Title 10 security cooperation* (CRS In Focus). Congressional Research Service.

Black, J. (2025, March 6). *David vs. Goliath: Cost asymmetry in warfare*. RAND Corporation. https://www.rand.org/pubs/commentary/2025/03/david-vs-goliath-cost-asymmetry-in-warfare.html

Center for Strategic and International Studies. (2024). *Empty bins in a wartime environment: The challenge of sustaining munitions stockpiles*. https://www.csis.org

Center for Strategic and International Studies. (2025). *From production lines to front lines: Assessing the U.S. defense industrial base*. https://www.csis.org

Center for Strategic and International Studies. (2026, February 17). *Small business, big impact: All about the base* [Event]. https://www.csis.org/events/small-business-big-impact-all-about-base

Clark, B., Patt, D., & Schadlow N. (2025). *Ending self-imposed scarcity: Exploiting America's commercial strengths to mobilize weapons production*. Hudson Institute. https://www.hudson.org

Commission on the National Defense Strategy. (2024). *Commission on the National Defense Strategy* (Report No. MSA-3057-4). RAND Corporation.

Congressional Research Service. (2020). *Security cooperation: An overview* (CRS In Focus No. IF11677). https://www.congress.gov/crs_external_products/IF/HTML/IF11677.web.html

Congressional Research Service. (2024). *Major Title 22 grant-based security assistance authorities and major Title 10 security cooperation authorities* (CRS Report No. R46337).

DoD. (2018). *DoD Instruction 3000.04: DoD munitions requirements process (MRP)* (Incorporating Change 2, August 31, 2018).

Fusco, S., Haney, K., Cartier, C. A., Olson, A., & Holmes, M. (2025). *Arsenal of policy: Defense industrial base wargame final report*. CNA.

GAO. (2002). *Defense management: Munitions requirements and combatant commanders' needs require linkage* (GAO-03-17). https://www.gao.gov/products/gao-03-17

GAO. (2024). *Critical materials: Action needed to implement requirements that reduce supply chain risks* (GAO-24-107176).

Greenwalt, W. C., & Patt, D. (2025). *Required to fail: Beyond documents—Accelerating joint advantage through direct resourcing and experimentation*. Hudson Institute.

Kane, T. M. (2016). *National security reform and the 2016 election* (CRS Report No. R44444). Congressional Research Service. https://sgp.fas.org/crs/natsec/R44444.pdf

Kress, A. J. (n.d.). *Proposed master DoD Munitions Requirements Process (MRP)* [Briefing document]. https://ia803206.us.archive.org/33/items/mipc-mccdc_usmc_mil/mccdc.usmc.mil_OTFGJ4KC3LJWJRLRTLJEI45TOKYYTJFD.pdf

Jones, S. G. (2023). *Empty bins in a wartime environment: The challenge to the U.S. defense industrial base*. Center for Strategic and International Studies. https://www.csis.org/analysis/empty-bins-wartime-environment-challenge-us-defense-industrial-base

Lehn, K. (2026, February). *Honolulu Defense Forum 2026 summary report*. Pacific Forum International.

Mitchell Institute for Aerospace Studies. (2023). *Building a force that wins: Munitions production and sustainment for great power conflict*. https://www.mitchellaerospacepower.org

Paparo, S. J. (2025, April 9). *Statement of Admiral Samuel J. Paparo, commander, U.S. Indo-Pacific Command: U.S. Indo-Pacific Command posture* [Testimony before the Senate Armed Services Committee]. U.S. Senate. https://www.armed-services.senate.gov/imo/media/doc/4102025fulltranscript.pdf

Paparo, S. J. (2026, January 12). *Remarks on the Indo-Pacific time–distance challenge* [Speech]. Honolulu Defense Forum, Honolulu, HI.

Sullivan, J. (2024, December 4). *Remarks by Assistant to the President for National Security Affairs Jake Sullivan on fortifying the U.S. defense industrial base* [Speech transcript]. The White House. https://bidenwhitehouse.archives.gov/briefing-room/speeches-remarks/2024/12/04/remarks-by-apnsa-jake-sullivan-on-fortifying-the-u-s-defense-industrial-base/

Taiwan Relations Act, Pub. L. No. 96-8, 93 Stat. 14 (1979).

United States Congress. (1986). *Goldwater–Nichols Department of Defense Reorganization Act of 1986* (Pub. L. No. 99-433). U.S. Government Printing Office. https://history.defense.gov/portals/70/documents/dod_reforms/goldwater-nicholsdodreordact1986.pdf

War on the Rocks. (2026, March). *Breadth—not just depth—is key to munitions industrial base resilience*. https://warontherocks.com/2026/03/magazine-breadth-not-just-depth-is-key-to-munitions-industrial-base-resilience/

Wasser, B., & Sheers, P. (2025, April). *From production lines to front lines: Revitalizing the U.S. defense industrial base for future great power conflict*. Center for a New American Security. https://www.cnas.org/publications/reports/from-production-lines-to-front-lines

THIS PAGE INTENTIONALLY LEFT BLANK

PANEL 16. ENGINEERING RESILIENCE IN CONTESTED LOGISTICS: RELAY-BASED HYBRID HUBS, TECHNICAL DATA OWNERSHIP, AND MANUFACTURING SIMULATION

Thursday, May 7, 2026, 1145 – 1300 ET (0845 - 1000 PT)

Panel Summary:

In the face of long distances, high-threat environments, and systemic data dependencies, traditional logistics models are no longer sufficient to ensure mission success. This panel explores cutting-edge strategies to secure the "Sustainment War" through three critical vectors: the construction of distributed logistics hubs on remote islands in the Western Pacific, the transformation of the Department into a Digital OEM to break vendor lock, and the use of advanced Modeling and Simulation (M&S) to link factory production capacity directly to field effectiveness. By integrating AI-powered grids with sea-based relay nodes, these researchers present a scalable framework for converting logistical vulnerabilities into an asymmetric warfighting advantage.

Chair: CDR Chris O'Connor, USN, Logistics Warfare Chair, Office of Warfare Studies, Naval Postgraduate School

Panel Presenters:

> **Constructing a Logistics Hub through the Integration of Remote Border Islands and Offshore Ocean Platforms –** *Captain Atsushi Yanagita, Japan Maritime Self-Defense Force*
>
> **Toward an AI-Powered Logistics Grid for Maintenance, Repair, and Overhaul: Why the DoW Must Become a Digital OEM to Win the Sustainment War –** *Jen Gebhardt, Govini*
>
> **From Factory to Field: Modeling Production Capacity and Logistics Effectiveness for Defense Systems –** *Michael Belisle, Northrop Grumman Corporation*

CDR Chris O'Connor, USN—is the Logistics Warfare Chair and Head of the Futures Group at the Naval Postgraduate School (NPS) Naval Warfare Studies Institute. A 2003 U.S. Naval Academy graduate, he specializes in strategic planning, force design, and contested logistics, previously serving at NATO Supreme Headquarters Allied Powers Europe (SHAPE).

Constructing a Logistics Hub through the Integration of Remote Border Islands and Offshore Ocean Platforms

Captain Atsushi Yanagita—is currently the Director of the Future Warfare and Logistics Studies Office at the Japan Maritime Self-Defense Force (JMSDF) Command and Staff College as well as a doctoral student in the Postgraduate School of Applied Environmental Systems at the Tokyo University of Marine Science and Technology. [y_atsushi1@inet.msdf.mod.go.jp, atsushi.yanagita@piano.ocn.ne.jp]

Abstract

This study examines whether a relay-based hybrid hub linking Japan's remote border islands with offshore ocean platforms can improve logistics performance in the Western Pacific under contested and disrupted conditions. Focusing on Minamitori Island and Okinotori Island, it addresses the sustainment gap between Japan, Guam, and Hawaii by comparing direct-route and relay-route models across six military and disaster-response scenarios. The concept is evaluated using four key performance indicators: Delivery, Distribution, Resilience, and Sustainment. Results indicate that relay routing consistently outperforms direct transport, especially in severe conditions. In the high-threat scenario, delivery probability rises from 47.9% to 65.6%, operational availability from 53.2% to 71.7%, and forward Days of Supply by about 2–3.5 days. These gains derive from a 72-hour buffer that absorbs disruption, supports limited inspection and repair, and enables redistribution in smaller lots. The study concludes that resilience through distribution provides a practical and scalable framework for sustaining maritime operations in contested environments.

Introduction

The principal Sea Lines of Communication (SLOCs) that support U.S. and allied maritime operations in the Indo-Pacific form a long and exposed arc extending from the continental United States through Hawaii and Guam to Japan. In the Western Pacific, this geography creates a persistent logistics problem: sustainment must be delivered across long distances, through harsh maritime conditions, and under the growing possibility of disruption by anti-access/area-denial (A2/AD) threats. The issue is therefore not simply how to replenish forces at sea, but how to preserve delivery reliability and sustainment endurance when routes are delayed, degraded, or contested.

At the center of this problem lies a specific operational gap. Between Japan, Guam, and Hawaii, there are few intermediate sustainment nodes capable of receiving, holding, redistributing, or recovering logistics flows. As a result, operational logistics remain heavily dependent on direct delivery and return-to-homeport cycles. This dependence reduces redundancy, constrains routing flexibility, and increases the likelihood that weather, port closures, mechanical interruptions, or hostile action will translate directly into a mission-level logistics failure.

Within this theater, Minamitori Island and Okinotori Island occupy positions of unusual strategic relevance. As remote Japanese border islands, they contribute to sovereignty, maritime domain awareness, and the maintenance of Japan's extensive Exclusive Economic Zone (EEZ). They also present a largely underdeveloped opportunity for distributed sustainment. Japanese policy has already supported incremental infrastructure development on remote islands, including berthing, mooring, and limited port functions. At the same time, Japan's National Defense Strategy and Defense Buildup Program, both issued in 2022, emphasize resilience, sustainment, automation, and manpower-efficient capability design. Taken together, these trends suggest that remote islands should be considered not only as symbolic or administrative outposts, but also as potential components of an operational logistics architecture.

Yet an important gap remains between policy direction and operational design. While current defense and infrastructure policies support resilience and remote-island development in general terms, they do not by themselves provide a concrete sustainment architecture for the mid-ocean logistics gap in the Western Pacific. Existing afloat logistics capabilities, including Combat Logistics Force operations, underway replenishment, and broader sea-basing concepts, remain indispensable. However, they continue to depend heavily on secure access to ports, airfields, anchorages, and fixed support infrastructure. In a disrupted or contested environment, those dependencies may become critical vulnerabilities rather than reliable enablers.

This study addresses that gap by proposing a relay-based hybrid logistics hub that integrates remote border islands with offshore ocean platforms. Rather than treating Minamitori Island and Okinotori Island as isolated points, the concept treats them as distributed sustainment nodes within a wider network linking Japan, Guam, Hawaii, and Okinawa. The core operational logic is straightforward: by inserting intermediate nodes capable of pre-positioning supplies, absorbing temporary disruption, enabling limited maintenance, and supporting onward redistribution, the logistics system can convert some forms of mission-ending delivery failure into manageable delay. In this sense, the concept is not designed primarily to maximize peacetime efficiency, but to preserve continuity under conditions of friction.

Because the maritime environment around both islands is characterized by typhoons, high sea states, and limited fixed infrastructure, the proposed architecture emphasizes mobility, modularity, and survivability. It therefore considers hybrid offshore support forms, including semi-submersible structures and Floating Production, Storage, and Offloading (FPSO)-derived configurations, as part of a scalable sustainment design. This approach is consistent with the logic of Distributed Maritime Operations (DMO), Expeditionary Advanced Base Operations (EABO), and contested logistics more broadly: combat power and sustainment capacity must be dispersed, adaptable, and able to recover under disruption rather than concentrated in a small number of fixed locations.

Against that background, this paper addresses two research questions. First, to what extent can a relay-based hybrid hub integrating remote border islands and offshore platforms improve logistics effectiveness in the Western Pacific under contested conditions? Second, which design features most strongly shape that effectiveness, particularly redundancy, rerouting capability, autonomy and uncrewed operations, and semi-submersible offshore storage?

To answer these questions, the paper evaluates the proposed concept through four key performance indicators (KPIs): Delivery, Distribution, Resilience, and Sustainment. Delivery refers to the probability that supplies successfully arrive at the forward destination. Distribution refers to the probability that at least one logistics route remains viable under disruption. Resilience refers to operational availability under degraded conditions. Sustainment refers to forward Days of Supply (DOS), or the duration for which operational demand can be covered at the forward edge. Together, these measures allow the analysis to assess not merely whether the concept is feasible in engineering terms, but whether it materially improves logistics performance under high-friction conditions.

This paper identifies the mid-ocean sustainment gap, proposes a relay-based hybrid hub linking remote islands and offshore platforms, and evaluates its effects on delivery assurance, routing robustness, operational continuity, and forward sustainment.

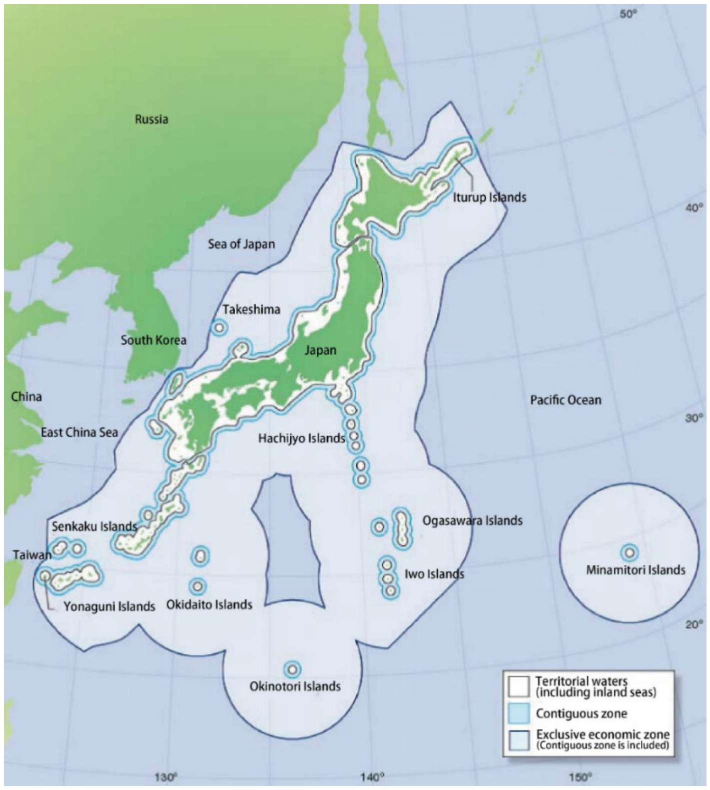

Figure 1. Japan's EEZ and Border Islands-

Problem Statement and Analytical Framework

Background and Problem Definition

Remote border islands are commonly understood as instruments of sovereignty, maritime security, and resource governance. In Japan's case, these functions are supported through a combination of legal authority, infrastructure investment, and administrative continuity. The Act on Special Measures for the Preservation of Remote Border Islands and related infrastructure programs have sustained habitation, transportation access, and basic port functions on selected islands. These measures are important not only for territorial administration, but also because they preserve the minimum physical and institutional conditions from which future sustainment functions may develop.

For the purposes of this study, however, the significance of remote islands is not primarily legal or symbolic. It is operational. In the Western Pacific, long sea distances, severe weather exposure, limited intermediate port capacity, and the growing possibility of disruption under anti-access/area-denial (A2/AD) conditions combine to undermine logistics reliability. Under these conditions, the central problem is not merely the movement of supplies from origin to destination, but the absence of intermediate nodes capable of receiving, holding, redistributing, and recovering logistics flows when direct delivery is delayed or interrupted.

Japan's recent defense policy sharpens this requirement. The National Defense Strategy, the Defense Buildup Program, and subsequent white papers place resilience, sustainment, automation, and manpower efficiency near the center of future force design. Read operationally, these documents imply two requirements. First, logistics systems must reduce excessive dependence on a small number of fixed bases. Second, they must become more capable of integrating civilian transportation capacity, modular infrastructure, and remote or low-manning support functions under appropriate legal and command arrangements. Yet policy direction alone does not resolve the theater-level architecture problem. The question remains how these principles should be translated into a practical sustainment network across the mid-ocean gap of the Western Pacific.

This paper defines that gap as a network problem. Existing logistics systems remain overly dependent on direct delivery and return-to-homeport cycles. Such dependence constrains routing options, limits redundancy, and makes disruption disproportionately costly. Weather, port denial, mechanical interruption, or hostile action can therefore quickly propagate from local friction to theater-level sustainment failure. The problem addressed here is how to reduce that fragility through a distributed architecture centered on relay-based hybrid hubs linking remote islands and offshore support forms.

Prior Research and Theoretical Lineage

The proposed concept is grounded in a long intellectual tradition in logistics and maritime operations. Classical logistics thought identified the enduring relationship between lines of operation, bases, and combat endurance. Jomini (1862) emphasized the strategic importance of lines of communication. Thorpe (1986) treated logistics as a national function re-quiring advanced preparation and coordination. Eccles (1959), writing from a naval perspective, argued that maritime operations can be sustained only through effective bases and supply systems. Across these traditions, a common proposition emerges: operational endurance depends not simply on the quantity of supply, but on the structure and reliability of the network that connects routes, nodes, and time.

Contemporary doctrine extends this logic rather than replacing it. Distributed Maritime Operations (DMO), Expeditionary Advanced Base Operations (EABO), sea-basing concepts, and the wider focus on contested logistics all point toward the same analytical conclusion: sustainment must be dispersed, adaptive, and resilient under conditions of disruption. In this respect, logistics architecture becomes a warfighting variable rather than a purely administrative function. The value of distribution lies not only in efficiency, but in preserving operational continuity when access is degraded, timing is uncertain, and attrition or interruption is expected.

At the same time, an important gap remains in the literature and in practical design. Existing research has not sufficiently integrated Japan's remote border islands—particularly Minamitori Island and Okinotori Island—with offshore ocean platforms as a combined sustainment architecture. Nor has prior work adequately linked such a concept to questions of phased implementation, acquisition feasibility, and measurable logistics performance. This study, therefore, builds on the theoretical lineage of logistics and current operational doctrine while addressing a specific applied problem that remains underdeveloped in the existing literature.

Research Questions (RQs) and KPIs

From this problem setting, two research questions follow.

- RQ-1: To what extent can a relay-based hybrid hub integrating remote border islands and offshore platforms improve logistics effectiveness in the Western Pacific under contested conditions?
- RQ-2: Which design features most strongly shape that effectiveness, particularly redundancy, rerouting capability, autonomy, and uncrewed operations, and semi-submersible off-shore storage?

These questions reflect a deliberate analytical sequence. The first question asks whether the proposed concept improves system performance relative to more direct, less distributed logistics arrangements. The second asks why such improvement occurs and which design attributes are most responsible for it. In other words, RQ-1 addresses comparative effectiveness, while RQ-2 addresses causal design significance.

The study's analytical logic is based on a network view of sustainment. The effectiveness of the proposed concept depends less on the absolute capacity of any single island or platform than on the network's ability to absorb disruption, preserve routing options, and maintain forward support over time. For that reason, the analysis treats modularity, dispersal, redundancy, and rapid re-tasking not as secondary engineering preferences, but as primary operational variables.

Key Performance Indicators and Evaluation Criteria

To evaluate the research questions, the study uses four key performance indicators (KPIs): Delivery, Distribution, Resilience, and Sustainment. In combination, these metrics capture both the immediate and cumulative effects of logistics architecture under contested conditions. The purpose is not merely to test whether supplies can move, but also to determine whether the system can continue functioning when exposed to friction, delay, and partial disruption.

These KPIs are analytically complementary. Delivery captures arrival success. Distribution captures route survivability. Resilience captures continuity under degradation. Sustainment captures temporal depth. Taken together, they provide a practical framework for comparing direct-delivery and relay-based logistics architectures in a theater where disruption is expected rather than exceptional.

Operational Constraints and Research Implications

The Western Pacific logistics problem is shaped by four interacting constraints: distance, weather, threat, and limited port capacity. These constraints do not operate independently. Instead, they reinforce one another. Long transit distances increase exposure time, severe sea states reduce schedule reliability, limited ports compress rerouting options, and threat conditions magnify the operational consequences of delay or concentration. The result is a logistics environment in which friction accumulates and is often nonlinear.

This has two implications for the present study. First, the relevant measure of effectiveness is not simply throughput under ideal conditions, but performance under disruption. Second, logistics resilience must be treated as a network property rather than as a feature of a single platform or facility. A contested logistics architecture succeeds when it can absorb delay, preserve at least partial functionality, and maintain forward momentum despite degraded conditions. For that reason, the analysis in the following sections emphasizes distributed sustainment nodes, pre-positioning, rerouting, modular support, and delay absorption as the principal mechanisms through which remote-island and offshore logistics hubs may improve theater sustainment.

Operational Environment and Geographic Conditions

Operational Concepts and Logistics Requirements

In the Indo-Pacific, logistics effectiveness is shaped by distance, exposure, and the ability to sustain forward forces amid uncertainty. Operational reach depends not only on lift capacity but also on whether supplies can be moved, held, redistributed, and recovered across long, vulnerable sea lines of communication (SLOCs). In this theater, contested logistics is therefore a problem of network performance rather than simple throughput. The key question is whether the sustainment system can preserve continuity when weather, sea-state conditions, port access limitations, mechanical interruption, or hostile action disrupt normal movement.

This requirement is consistent with the operational logic of Distributed Maritime Operations (DMO) and Expeditionary Advanced Base Operations (EABO). Both concepts seek to dis-perse combat power and supporting functions in ways that complicate adversary targeting while preserving mission effectiveness. From a logistics perspective, this implies more than additional shipping or larger inventories. It requires a distributed sustainment architecture composed of intermediate nodes, pre-positioned stocks, flexible command-and-control (C2), and the ability to reprioritize and reroute support under degraded conditions. In other words, sustainment must be designed not only for efficiency in routine conditions, but for recovery and continuity under friction.

The Western Pacific places unusual pressure on this requirement. Transit distances are long, exposure times are extended, and the number of usable intermediate logistics locations remains limited. As a result, direct delivery often becomes the default operating model. That model can function in permissive conditions, but it becomes increasingly brittle when even modest disruption accumulates across multiple variables. The operational value of an intermediate relay node is therefore not simply geographic convenience. It lies in the ability to convert disruption from mission termination into bounded delay, while preserving options for onward movement and local recovery. This is the operational requirement against which Minamitori Island and Okinotori Island are assessed in the present study.

Geographic Characteristics of Minamitori Island and Okinotori Island

Minamitori Island offers comparatively favorable conditions for the development of an initial logistics node. Existing facilities, anchorage potential, and limited air and sea access make it more suitable as a first-step site for phased logistics expansion. Relative to more austere offshore locations, it presents a more practical environment for introducing pre-positioned stocks, transshipment functions, light maintenance support, and command-and-control connectivity. In analytical terms, Minamitori Island is significant because it can serve as a feasible entry point for a distributed sustainment architecture rather than requiring a fully mature offshore system from the outset. Okinotori Island presents a markedly different operating environment. It is more austere, more environmentally demanding, and less suited to conventional fixed-facility development. These conditions increase the relative value of offshore, modular, semi-submersible, and FPSO-derived support forms. In this setting, the central engineering problem is not how to reproduce a conventional port, but how to provide logistics persistence under severe maritime conditions with limited permanent infrastructure. The importance of Okinotori Island, therefore, lies less in fixed basing potential than in its role as a location where survivable and relocatable offshore support concepts become operationally relevant.

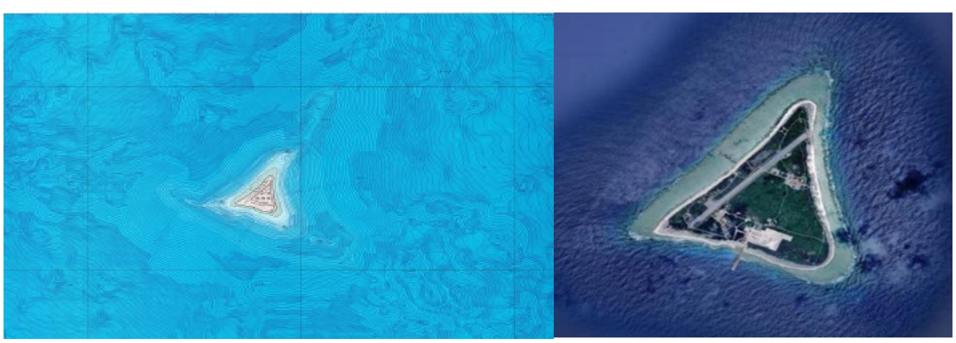

Figure 2. Minamitori Island

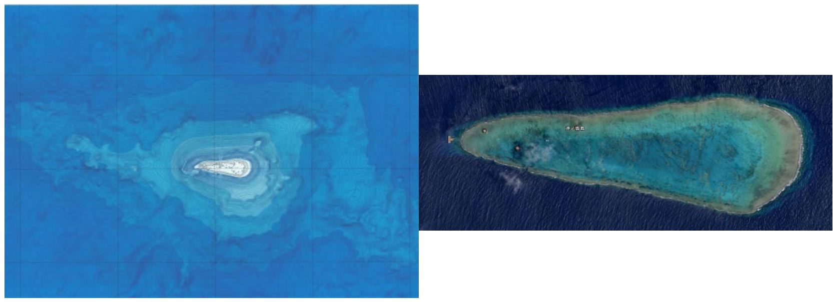

Figure 3. Okinotori Island

The contrast between the two islands is analytically useful. Minamitori Island represents a comparatively accessible and developmentally realistic node for early implementation. Okinotori Island represents the more demanding edge case that tests the necessity of modularity, mobility, and survivability. Taken together, they define a scalable design spectrum for West-ern Pacific logistics: one anchored in practical phased development, and the other in engineering adaptation to austere and contested maritime conditions. This duality is central to the concept developed in the following chapter.

Both islands are also significant at the theater level. Beyond their sovereignty and Exclusive Economic Zone (EEZ) functions, they may serve as intermediate connectors within a wider allied and partner logistics network linking Japan, Guam, Hawaii, and forward operating areas such as Okinawa. Their value is therefore not reducible to local infrastructure alone. It emerges from their position within a broader network of routes, stocks, and relay functions across the Western Pacific.

Strategic Significance as a Western Pacific Logistics Hub

The strategic significance of a logistics hub centered on Minamitori Island and Okinotori Is-land is best understood through the four key performance indicators used in this study: De-livery, Distribution, Resilience, and Sustainment. The value of the concept lies not in any single engineering feature, but in the way geographic position and hybrid-node design affect all four dimensions of logistics performance.

First, Delivery may improve because the islands and their associated offshore platforms can function as forward replenishment and redistribution nodes. By inserting an intermediate node between origin and destination, the network reduces exclusive reliance on uninterrupted direct delivery. This enables greater ability to absorb delays, stage critical supplies in advance, and shorten the effective response cycle for onward movement to the forward edge.

Second, Distribution may improve because the combination of island facilities with floating, semi-submersible, and FPSO-derived platforms increases routing flexibility. A hybrid architecture creates multiple support forms rather than a single point of dependence. In a disrupted environment, that redundancy matters because the probability of maintaining at least one viable route or node is often more important than the nominal capacity of any one facility.

Third, Resilience may improve because semi-submersible and modular offshore support structures are potentially more tolerant of severe weather, partial damage, and access degradation than conventional fixed-port assumptions would suggest. In an environment where disruption is expected, the ability to maintain partial functionality and recover operational availability becomes a core source of advantage. For this reason, survivability should be treated not only as a platform attribute, but as a logistics-network attribute.

Fourth, Sustainment may improve because pre-positioned stocks and local replenishment cycles at these nodes increase forward Days of Supply (DOS). In practical terms, this ex-tends the time forward forces can operate without immediate reliance on long-haul resupply from rear bases. The concept, therefore, improves not only movement reliability but also endurance at the point of operational need.

Taken together, these implications show why the geographic characteristics of Minamitori Island and Okinotori Island matter strategically. Their significance does not arise merely from location on a map. It arises from the combination of geographic position, infrastructure potential, offshore engineering adaptability, and network function. In that sense, they should be understood not as isolated outposts, but as candidate relay nodes in a broader theater sustainment architecture.

Architecture of the Logistics Hub

Design Philosophy and Minimum Viable Product (MVP)

The proposed logistics hub should be understood not as a single fixed installation, but as a distributed, scalable architecture that combines remote-island infrastructure with offshore support systems. Its central design logic is evolutionary rather than static. Instead of seeking to construct a fully mature offshore base from the outset, the concept begins with a minimum viable architecture capable of producing operationally meaningful improvement in logistics performance under disrupted conditions and then expands in phases as legal arrangements, technical confidence, and mission demand mature. This design philosophy is consistent with contested logistics, where early functionality, adaptability, and survivability are often more valuable than initial completeness.

For the purposes of this study, the minimum viable product (MVP) is defined not by physical size, but by functional sufficiency. A node qualifies as operationally meaningful when it can do more than merely hold cargo. At a minimum, the initial architecture must be able to receive supplies, hold them safely for short periods, redistribute them onward, support limited inspection and light recovery functions, and maintain enough communications and monitoring capacity to operate remotely under degraded conditions. In practical terms, the MVP serves as the first relay point in a broader sustainment network. Its purpose is not to replicate a conventional port, but to create a node that can absorb disruption and preserve continuity.

This framing has two important implications. First, the architecture is optimized initially for operational usefulness rather than permanence. Second, capability growth should be evaluated by measurable gains in Delivery, Distribution, Resilience, and Sustainment rather than by construction scale alone. In this sense, the logistics hub is best understood as a performance-based architecture whose value depends on what functions it enables across the network, not simply on the amount of infrastructure emplaced at a single location.

Structural Portfolio and Operational Fit

The proposed hub can be realized through multiple structural forms, each contributing differently to logistics performance. The central architectural question is therefore not which single structure is best, but how a portfolio of structures can be combined so that each offsets the others' limitations while contributing to the four key performance indicators. The relevant structural categories in this study are fixed island or pier-type facilities, floating barges and modular pontoons, semi-submersible platforms, and FPSO-derived structures such as con-verted tankers.

Taken together, these structural options indicate that the logistics hub is best treated as a portfolio design problem. Minamitori Island is better suited to phased development that begins with fixed and modular floating elements, whereas Okinotori Island is more naturally aligned with offshore-first configurations centered on semi-submersible and FPSO-derived support forms. This distinction follows directly from the differences in access, environmental severity, and infrastructure potential identified in the previous chapter.

Table 1. Structural Typology (Synopsis)

Structure Type	***Primary Functional Significance***	***KPI(s)***
Fixed: artificial island/pier 	Fixed structures provide continuity, visibility, and political legitimacy. They are especially useful where stable access, routine transfer, and a visible sovereign presence are required. Their principal advantage lies in delivery assurance under relatively controlled conditions. Their principal weakness is their least adaptability and potential vulnerability to both adversary targeting and environmental stress. For that reason, they should be treated as one element of the architecture rather than its complete solution.	Delivery
Floating: barge / modular pontoon 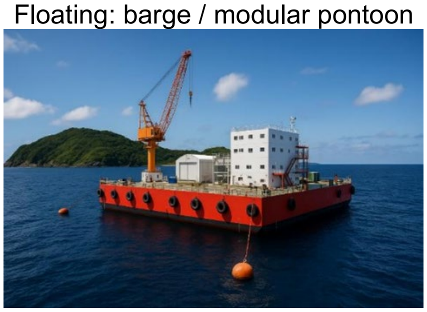	Floating barges and modular pontoons provide rapid fielding, flexible augmentation, and lower barriers to modification. These attributes make them attractive in early implementation phases, when rapid delivery capacity and basic distribution flexibility are more important than full survivability. They are especially useful for establishing an initial logistics presence at a lower cost and with shorter lead times. Their limitations become more apparent in severe maritime conditions, where endurance and survivability are weaker than those of more specialized offshore forms.	Delivery / Distribution
Semi-submersible 	Semi-submersible platforms become more important as weather survivability and operational continuity become central design drivers. Their relative stability in high sea states, lower visibility profile, and greater tolerance of severe maritime conditions make them particularly relevant to resilience. In austere environments such as Okinotori Island, where a conventional fixed-facility solution is less feasible, semi-submersible forms offer a practical way to preserve logistics functionality when simpler floating structures may be forced to suspend operations.	Resilience / Sustainment

FPSO-derived (converted tanker) 	FPSO-derived structures add bulk storage and sustained offshore support capacity. Their main value lies in fuel, water, and critical materiel storage, as well as in their capacity to support repeated redistribution cycles. They therefore strengthen sustainment depth and distribution endurance more than political visibility or symbolic presence. Within the larger portfolio, they serve as the backbone of offshore sustainment rather than as the entire hub concept on their own.	Sustainment / Distribution

Required Functions and Their KPI Relevance

From an operational standpoint, the architecture must perform four core functions: access, storage, maintenance, and management. These categories are not merely engineering descriptors. They are the practical mechanisms through which the hub improves logistics performance under contested or disrupted conditions.

- **Access** includes berthing and mooring, fuel and water transfer, container handling, and lighterage. These functions are foundational because no relay architecture can work un-less cargo can be received and moved onward reliably. Access, therefore, maps most directly to Delivery.
- **Storage** includes automated warehousing, unmanned fuel tanks, and containerized stockholding. Storage matters because it decouples inbound and outbound movement. In other words, it allows the hub to absorb temporary disruption rather than forcing immediate mission failure when onward movement is delayed. Storage therefore strengthens both Distribution and Sustainment.
- **Maintenance** includes recovery, inspection, charging, relaunch, and limited repair capability for uncrewed systems and selected equipment. These functions reduce effective down-time, restore degraded capability, and shorten turnaround time. Maintenance, therefore, contributes directly to Resilience.
- **Management** includes communications control, medical support, berthing coordination, and power generation. Communications are particularly critical because the hub must operate under remote, distributed, and sometimes unmanned conditions. A multi-path communications architecture—combining satellite, HF, optical, and other available means—improves continuity under degraded conditions. Similarly, hybrid energy arrangements strengthen endurance in routine operations while preserving resilience during crisis or severe weather. Management, therefore, is the integrating function that allows the node to behave as an active part of a logistics network rather than as a passive storage point.

Phased Capability Development

Because legal authority, technical maturity, operational concepts, and contracting mechanisms will not mature at the same pace, the proposed architecture should be developed in phases. The phased growth model in this study begins with an initial fielded MVP largely based on civilian technology, with an emphasis on legal compliance, safety, and a rapid-delivery baseline. It then expands to include fuel, water, and maintenance functions, followed by larger warehousing and management capacity, and finally advanced communications and autonomous control that support performance across all four KPIs.

This should not be interpreted as a simple equipment sequence. It is better understood as a spiral development process in which operational employment, policy adaptation, safety regulation, engineering choices, and data infrastructure mature together. That point is especially important from an acquisition perspective. The challenge is not merely to procure a set of platforms, but to build an integrated sustainment system whose legal, technical, and operational components become progressively more capable and mutually reinforcing over time.

Phased development also provides a practical response to uncertainty. It allows the architecture to begin with achievable capabilities, generate early operational utility, and then adapt based on evidence about which combinations of structures and functions produce the greatest improvement in Delivery, Distribution, Resilience, and Sustainment. In this respect, phased growth is not only an implementation strategy. It is also an analytical strategy that supports iterative learning under real-world constraints.

Architectural Significance

The central implication of this chapter is that the proposed logistics hub should be understood as a distributed, layered, and adaptive sustainment system rather than as a single offshore installation. Fixed infrastructure provides continuity and legitimacy. Floating modular elements provide rapid fielding and flexible augmentation. Semi-submersible structures provide survivability under severe environmental and threat conditions. FPSO-derived assets provide sustainment depth through bulk offshore storage and the capacity for repeated redistribution. The value of the architecture does not lie in any one of these forms alone, but in the way they are combined to generate network effects greater than those of any single structure.

This design logic establishes the basis for the analytical chapter that follows. If the architecture is valid, its effects should be visible not only in engineering feasibility but also in measurable changes to logistics performance across different operational scenarios. The next chap-ter, therefore, evaluates whether the proposed relay-based hybrid hub actually changes the dominant failure mode of logistics operations from mission-ending non-delivery to bounded and operationally manageable delay, and whether that shift yields meaningful gains across the four key performance indicators.

Scenarios and Evaluation

Analytical Purpose and Comparative Design

This chapter evaluates whether a relay-based hybrid hub improves logistics performance in the Western Pacific relative to a direct-route baseline. The comparison is made between two routing constructs. The first is a direct-route model, in which cargo moves from the origin to the destination without passing through an intermediate node. The second is a relay-route model, in which Minamitori Island and Okinotori Island function as intermediate distributed sustainment nodes (DSNs) that absorb disruption, support limited maintenance, and enable onward redistribution. The analytical objective is to determine whether the relay construct shifts the dominant failure mode of logistics operations from mission-ending non-delivery to bounded delay that can be managed operationally.

To test that proposition, the analysis uses two scenario families. The military scenarios (M-series) represent progressively more contested operating environments, while the non-military scenarios (C-series) represent escalating disruption in civil and disaster-response conditions. Across both families, the comparison remains the same: whether relay routing preserves delivery, resilience, and sustainment more effectively than direct transport as operational friction increases.

Measures of Effectiveness

The concept is evaluated through four key performance indicators (KPIs): Delivery, Distribution, Resilience, and Sustainment. Delivery is defined as the probability that supplies successfully arrive at Okinawa. Distribution is defined as the probability that at least one logistics route remains viable under disruption. Resilience is defined as operational availability under degraded conditions. Sustainment is defined as forward Days of Supply (DOS), that is, how long operational demand at the forward destination can be covered. Together, these KPIs enable evaluation of the proposed hub not merely as infrastructure, but as a comparative sustainment architecture under contested and disrupted conditions.

Scenario Families

To preserve internal consistency across the six scenarios, the comparative model uses a common set of scenario-specific representative assumptions. These values are not presented as universal constants or fleet-wide empirical rates. Rather, they are structured analytical inputs used to generate the comparative outputs reported in Table 3.

Weather disruption probability is set at 0.03 / 0.03 / 0.10 for M-1/M-2/M-3 and 0.03 / 0.15 / 0.15 for C-1/C-2/C-3.

Port disruption probability is set at 0.02 / 0.10 / 0.30 for M-1/M-2/M-3 and 0.02 / 0.20 / 0.15 for C-1/C-2/C-3. The share of disruptions exceeding the 72-hour holding buffer is set at 0.10 / 0.20 / 0.40 for M-1/M-2/M-3 and 0.10 / 0.30 / 0.50 for C-1/C-2/C-3.

Mechanical failure probability varies by routing construct: for military scenarios, direct values are 0.0198 / 0.0257 / 0.0488, and relay values are 0.0150 / 0.0180 / 0.0300; for non-military scenarios, direct values are 0.0060 / 0.0072 / 0.0084, and relay values are 0.0050 / 0.0065 / 0.0075.

Threat-related disruption is applied only in military scenarios and is set at 0.055 / 0.090 / 0.200 for M-1/M-2/M-3, whereas in non-military scenarios, the threat is set to zero.

These assumptions reflect the paper's routing logic. In the direct case, weather, port closures, mechanical interruptions, and threats are treated as immediate causes of mission failure. In the relay case, only the share of weather- and port-related disruptions that exceed the 72-hour hold window remains mission-ending, while inspection and light repair at the intermediate node reduce the effective mechanical-interruption rate.

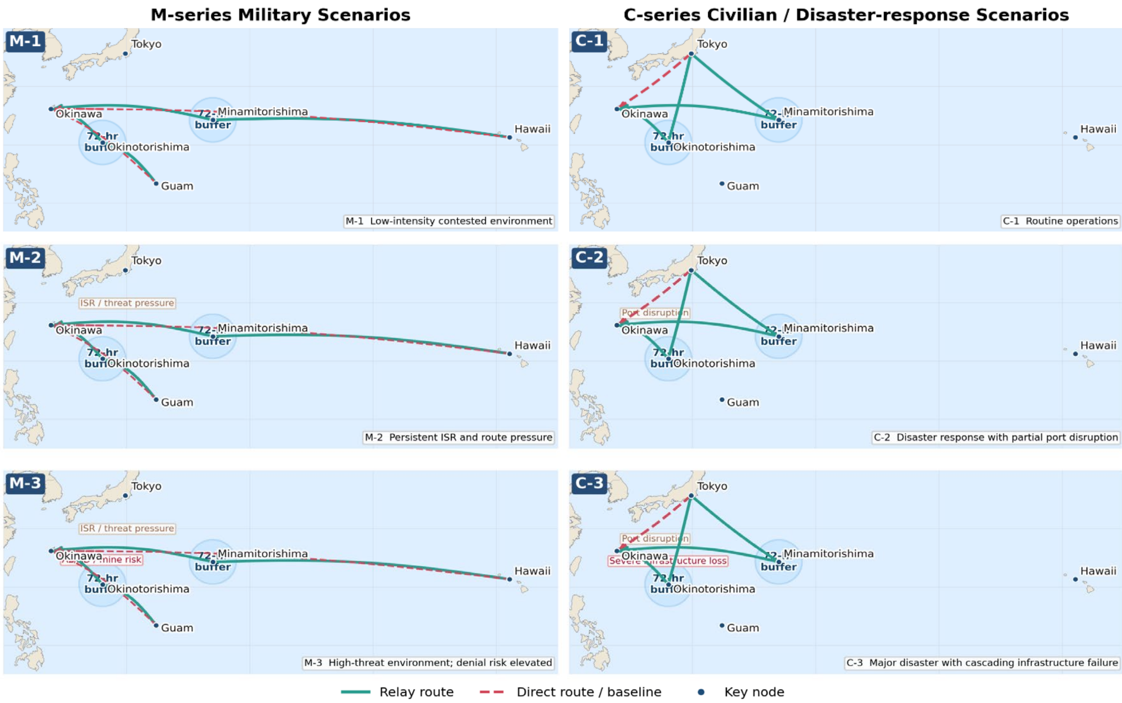

Figure 4. Military Scenario (M-Series) and Non-Military Scenario (C-Series)

Comparative Calculation Logic

Delivery

Delivery is modeled as the probability that supplies successfully reach Okinawa. For the direct-route construct in scenarios, arrival probability is defined as

$$D_s^{\text{direct}} = (1 - W_s)(1 - M_s^{\text{direct}})(1 - P_s)(1 - T_s) \qquad (1)$$

where W_s is weather disruption probability, M_s^{direct} is direct-case mechanical failure probability, P_s is port disruption probability, and T_s is threat-related disruption probability. This is a series-reliability formulation: if any one of these factors causes mission non-completion, delivery fails.

For the relay-route construct, the same baseline risk categories apply, but the 72-hour holding buffer changes the treatment of weather and port effects, and relay-node inspection/light repair reduces the mechanical interruption rate. Arrival probability is therefore defined as

$$D_s^{\text{relay}} = (1 - \alpha_s W_s)(1 - M_s^{\text{relay}})(1 - \alpha_s P_s)(1 - T_s) \qquad (2)$$

where α_s is the share of weather- and port-related disruptions exceeding the 72-hour hold window and M_s^{relay} is relay-case mechanical failure probability after inspection and light repair at the intermediate node. For non-military scenarios, $T_s = 0$. This formulation captures the central mechanism of the relay concept: part of the disruption is no longer counted as immediate non-delivery.

Using the representative assumptions above, the model reproduces the Delivery values shown in Table 3.

For M-1:

$$D_{M1}^{\text{direct}} = (1 - 0.03)(1 - 0.0198)(1 - 0.02)(1 - 0.055) = 0.881$$
$$D_{M1}^{\text{relay}} = (1 - 0.10 \times 0.03)(1 - 0.0150)(1 - 0.10 \times 0.02)(1 - 0.055) = 0.92$$

For M-2:

$$D_{M2}^{\text{direct}} = (1 - 0.03)(1 - 0.0257)(1 - 0.10)(1 - 0.090) = 0.774$$
$$D_{M2}^{\text{relay}} = (1 - 0.20 \times 0.03)(1 - 0.0180)(1 - 0.20 \times 0.10)(1 - 0.090) = 0.870$$

For M-3:

$$D_{M3}^{\text{direct}} = (1 - 0.10)(1 - 0.0488)(1 - 0.30)(1 - 0.200) = 0.479$$
$$D_{M3}^{\text{relay}} = (1 - 0.40 \times 0.10)(1 - 0.0300)(1 - 0.40 \times 0.30)(1 - 0.200) = 0.656$$

For non-military scenarios, because threat is zero, the equations simplify to

$$D_s^{\text{direct}} = (1 - W_s)(1 - M_s^{\text{direct}})(1 - P_s) \qquad (3)$$
$$D_s^{\text{relay}} = (1 - \alpha_s W_s)(1 - M_s^{\text{relay}})(1 - \alpha_s P_s) \qquad (4)$$

Thus, for C-1:

$$D_{C1}^{\text{direct}} = (1 - 0.03)(1 - 0.0060)(1 - 0.02) = 0.945$$
$$D_{C1}^{\text{relay}} = (1 - 0.10 \times 0.03)(1 - 0.0050)(1 - 0.10 \times 0.02) = 0.990$$

For C-2:

$$D_{C2}^{\text{direct}} = (1 - 0.15)(1 - 0.0072)(1 - 0.20) = 0.675$$
$$D_{C2}^{\text{relay}} = (1 - 0.30 \times 0.15)(1 - 0.0065)(1 - 0.30 \times 0.20) = 0.892$$

For C-3:

$$D_{C3}^{\text{direct}} = (1 - 0.15)(1 - 0.0084)(1 - 0.15) = 0.716$$
$$D_{C3}^{\text{relay}} = (1 - 0.50 \times 0.15)(1 - 0.0075)(1 - 0.50 \times 0.15) = 0.849$$

Distribution

Distribution is modeled as the probability that at least one logistics route remains viable under disruption. Conceptually, this corresponds to a parallel-system reliability problem:

$$G_s = 1 - \prod_{r=1}^{n} \left(1 - \beta_s p_{r,s}\right) \qquad (5)$$

where p_(r,s)denotes the single-route arrival probability for routerin scenario s and β_sis an effective independence factor used to account for common-cause risk such as shared weather or regional disruption. In the military scenarios, the route set consists of the Hawaii line and the Guam line. In non-military scenarios, the route set comprises the direct Tokyo–Okinawa path and two relay options via Minamitori Island and Okinotori Island.

The model uses Distribution primarily as a redundancy indicator rather than as the central headline result, but its logic supports the broader claim that relay routing preserves at least one viable path under stress.

Resilience

Resilience is defined as operational availability under degraded conditions. In this study, resilience is modeled as the product of mechanical availability and operational availability:

$$R_s = A_s^{\text{mech}} \cdot A_s^{\text{op}} \qquad (6)$$

where R_s is resilience in scenario s, A_s^{mech} is mechanical availability, and A_s^{op} is operational availability.

For the representative setting used in this study, mechanical availability is set at

$$A_s^{\text{mech,direct}} = 1 - M_s^{\text{direct}} \qquad (7)$$
$$A_s^{\text{mech,relay}} = 1 - M_s^{\text{relay}} \qquad (8)$$

The direct value represents baseline mechanical availability, while the relay value represents improved effective availability after inspection and light repair at the intermediate node.

Operational availability is defined as the complement of effective downtime. For the direct-route construct,

$$A_s^{\text{op,direct}} = 1 - L_s^{\text{direct}} \qquad (9)$$

$$L_s^{\text{direct}} = W_s + P_s + T_s \qquad (10)$$

Here, W_s denotes weather-related downtime, P_s port-related downtime, and T_s threat-related stoppage in scenarios.

For the relay-route construct, the 72-hour holding buffer reduces the weather- and port-related share of effective downtime so that

$$A_s^{\text{op,relay}} = 1 - L_s^{\text{relay}} \qquad (11)$$

$$L_s^{\text{relay}} = \alpha_s W_s + \alpha_s P_s + T_s \qquad (12)$$

where α_s denotes the share of weather- and port-related disruptions exceeding the 72-hour holding buffer. In other words, only the portion of weather and port disruption that cannot be absorbed within the holding window is counted as mission-degrading downtime in the relay case. This is the principal resilience mechanism of the relay-based hub.

Substituting these definitions gives

$$R_s^{\text{direct}} = (1 - M_s^{\text{direct}})\,[1 - (W_s + P_s + T_s)] \qquad (13)$$
$$R_s^{\text{relay}} = (1 - M_s^{\text{relay}})\,[1 - (\alpha_s W_s + \alpha_s P_s + T_s)] \qquad (14)$$

Using the representative mechanical availability values above, the operational availability implied by the reported resilience outcomes can be calculated directly.

For example, in M-1:

$$R_{M1}^{\text{direct}} = (1 - 0.0198)\,[1 - (0.03 + 0.02 + 0.055)] = 0.9344$$
$$R_{M1}^{\text{relay}} = (1 - 0.0150)\,[1 - (0.10 \times 0.03 + 0.10 \times 0.02 + 0.055)] = 0.9831$$

Likewise, in M-3:

$$R_{M3}^{\text{direct}} = (1 - 0.0488)\,[1 - (0.10 + 0.30 + 0.20)] = 0.5317$$
$$R_{M3}^{\text{relay}} = (1 - 0.0300)\,[1 - (0.40 \times 0.10 + 0.40 \times 0.30 + 0.20)] = 0.7166$$

The same logic applies to the non-military scenarios. In C-2, for example:

$$R_{C2}^{\text{direct}} = (1 - 0.0072)\,[1 - (0.15 + 0.20)] = 0.6751$$
$$R_{C2}^{\text{relay}} = (1 - 0.0065)\,[1 - (0.30 \times 0.15 + 0.30 \times 0.20)] = 0.8959$$

Across all scenarios, relay routing improves resilience by reducing effective downtime attributable to weather and port disruptions within the 72-hour hold window while also increasing mechanical availability through inspection and light repair at the intermediate node. Accordingly, the resilience results reported in Table 3 are interpreted not as a separate phenomenon from delivery, but as a related measure of how the relay construct improves operational continuity under stress.

Sustainment

Sustainment is measured as forward Days of Supply (DOS). In the model, DOS depends on forward stock, cargo deliverable per resupply cycle, arrival probability, and daily demand. Direct routing is represented as

$$DOS_s^{\text{direct}} = \frac{F + Q \cdot D_s^{\text{direct}}}{q} \qquad (15)$$

and relay routing as

$$DOS_s^{\text{relay}} = \frac{F + C + Q \cdot D_s^{\text{relay}}}{q} \qquad (16)$$

where $F$is forward stock, $Q$is per-cycle flow, $C$is relay-node consolidated stock, $D$is Delivery probability, and $q$is daily demand. The assumed values are daily demand $q = 1{,}000$t, forward stock $F = 10{,}000$t, per-cycle flow $Q = 5{,}000$t, and relay-node consolidated stock $C = 3{,}000$t.

Dividing through by daily demand yields a convenient days-based form:

$$DOS_s^{\text{direct}} = 10 + 5D_s^{\text{direct}} \qquad (17)$$
$$DOS_s^{\text{relay}} = 13 + 5D_s^{\text{relay}} \qquad (18)$$

Thus, for M-1:

$$DOS_{M1}^{\text{direct}} = 10 + 5(0.881) = 14.4$$
$$DOS_{M1}^{\text{relay}} = 13 + 5(0.926) = 17.4$$

For M-2:

$$DOS_{M2}^{\text{direct}} = 10 + 5(0.774) = 13.9$$
$$DOS_{M2}^{\text{relay}} = 13 + 5(0.870) = 17.0$$

For M-3:

$$DOS_{M3}^{\text{direct}} = 10 + 5(0.479) = 12.4$$
$$DOS_{M3}^{\text{relay}} = 13 + 5(0.656) = 15.2$$

For C-1:

$$DOS_{C1}^{\text{direct}} = 10 + 5(0.945) = 14.8$$
$$DOS_{C1}^{\text{relay}} = 13 + 5(0.990) = 17.9$$

For C-2:

$$DOS_{C2}^{\text{direct}} = 10 + 5(0.675) = 13.4$$
$$DOS_{C2}^{\text{relay}} = 13 + 5(0.892) = 17.1$$

For C-3:

$$DOS_{C3}^{\text{direct}} = 10 + 5(0.716) = 13.6$$
$$DOS_{C3}^{\text{relay}} = 13 + 5(0.849) = 16.8$$

Results Overview

Across all six scenarios, the relay-route construct outperforms the direct-route construct on the principal measures of effectiveness. In the military series, Delivery improves from 88.1% to 92.6% in M-1, from 77.4% to 87.0% in M-2, and from 47.9% to 65.6% in M-3. Resilience improves from 93.44% to 98.31%, 81.49% to 92.17%, and 53.17% to 71.66%, respectively. DOS increases from 14.4 to 17.4 days, from 13.9 to 17.0 days, and from 12.4 to 15.2 days across the same scenarios.

A similar pattern appears in the non-military series. Delivery improves from 94.5% to 99.0% in C-1, from 67.5% to 89.2% in C-2, and from 71.6% to 84.9% in C-3. Resilience improves from 94.38% to 99.30%, 67.51% to 89.59%, and 71.73% to 85.39%, respectively. DOS rises from 14.8 to 17.9 days, 13.4 to 17.1 days, and 13.6 to 16.8 days. The largest delivery uplift appears in C-2, while the largest military uplift appears in M-3.

Table 2. KPI Summary by Scenario

Scenario	Delivery Direct	Delivery Relay	Δ Delivery	Distribution*	Resilience Direct	Resilience Relay	Δ Resilience	DOS Direct	DOS Relay	Δ DOS
M-1	88.1%	92.6%	+4.5 pts	97.3%	93.44%	98.31%	+4.87 pts	14.4	17.4	+3.0
M-2	77.4%	87.0%	+9.6 pts	95.0%	81.49%	92.17%	+10.68 pts	13.9	17.0	+3.1
M-3	47.9%	65.6%	+17.7 pts	77.4%	53.17%	71.66%	+18.49 pts	12.4	15.2	+2.8
C-1	94.5%	99.0%	+4.5 pts	99.99%	94.38%	99.30%	+4.92 pts	14.8	17.9	+3.1
C-2	67.5%	89.2%	+21.7 pts	98.6%	67.51%	89.59%	+22.08 pts	13.4	17.1	+3.7
C-3	71.6%	84.9%	+13.3 pts	98.3%	71.73%	85.39%	+13.66 pts	13.6	16.8	+3.2

Note. Delivery measures the probability of arrival at Okinawa. Distribution measures the probability that at least one logistics route remains viable. Resilience measures operational availability under degraded conditions. DOS indicates forward Days of Supply. Positive deltas indicate improvement under the relay-route construct.

Interpretation

The relay construct performs better for three reasons. First, it changes the temporal structure of disruption. In the direct model, weather and port closure often terminate the mission. In the relay model, only the share of disruption exceeding the 72-hour hold window remains mission-ending. Equations (1) through (4) make that mechanism explicit.

Second, the relay node improves operational availability through maintenance refresh. Inspection and light repair reduce the effective mechanical interruption rate and raise mechanical availability from 0.9929 to 0.9980 in the representative setting. Equations (6) through (14) show that this effect compounds with the reduction in weather- and port-related downtime.

Third, the relay architecture increases the forward time margin. In days-based form, direct routing produces DOS_s^"direct" =10+5D_s^"direct" , while relay routing produces DOS_s^"relay" =13+5D_s^"relay" . The additional three days are generated by relay-node consolidated stock and short-cycle shuttle resupply.

Taken together, these results show that the proposed relay-based hub is most valuable not in routine conditions, but in precisely those circumstances where direct-route logistics architectures are most likely to fail. The relay construct does not merely optimize throughput in peacetime; it changes how the logistics network behaves under stress.

Operational Significance

The operational meaning of these findings is straightforward. The proposed hub should not be understood primarily as a storage site, nor merely as an engineering installation. It is better understood as an operational control point that integrates time, space, and information. As a time mechanism, it creates a bounded window for delay absorption. As a space mechanism, it enables offshore handling and redistribution without full dependence on fixed-port access. As an information mechanism, it supports reprioritization, re-sorting, and near-real-time mission allocation under degraded conditions. This interpretation is consistent with the study's broader architecture and phased fielding logic.

Conclusion

This study examined whether a relay-based hybrid hub, using Japan's remote border islands and ocean platforms, can improve logistics performance in the Western Pacific under both contested and disrupted conditions. The analysis compared a direct-route model with a relay-route model centered on distributed sustainment nodes at Minamitori Island and Okinotori Island. Across all scenarios, the relay constructs improved Delivery, Distribution, Resilience, and Sustainment, with the largest gains appearing in the most severe military and disaster-response conditions. These findings indicate that the principal value of the proposed architecture lies not in increasing efficiency under routine conditions, but in changing the dominant failure mode of logistics operations. More specifically, the relay-based hub converts mission-ending delivery failure into a manageable delay through a 72-hour buffer, offshore handling capacity, and a shared logistics picture that supports rerouting and recovery.

The operational implication is clear. In a theater characterized by long distances, contested sea lines, infrastructure fragility, and uncertain access, logistics resilience must be designed as a network function rather than treated as a by-product of transportation capacity alone. The proposed hub demonstrates that time, space, and information can be integrated into a distributed sustainment architecture that preserves continuity even when direct delivery becomes unreliable. In this sense, remote islands are not merely geographic outposts. Properly integrated with maritime platforms and data-driven logistics control, they can function as operational shock absorbers within a broader sustainment network.

At the same time, this study should be understood as a first step rather than a final design solution. The present model demonstrates the utility of the relay concept, but further research is required to determine the most effective physical and operational form of the hub. One particularly promising direction is the application of Mega-Float technology, an area in which Japan possesses significant engineering experience and comparative advantage. If a Mega-Float structure were configured as a maritime logistics hub, it could provide a more flexible and effective base for offshore storage, transshipment, light maintenance, and distributed support functions in areas where fixed-port infrastructure is unavailable, vulnerable, or politically constrained.

Such a development would extend the logic of this paper in an important way. Whereas the current study establishes the value of relay-based logistics through remote islands and ocean platforms, a Mega-Float-based hub could strengthen that architecture by adding scalable sea-based infrastructure with greater modularity, persistence, and operational adaptability. It may also improve the hub's ability to support multi-domain logistics missions, including sustainment, repair, medical support, and temporary command-and-control functions. In this respect, Mega-Float is not simply an engineering add-on, but a possible next-generation embodiment of the relay-based hybrid hub concept.

Accordingly, future research will examine how Mega-Float technology can be incorporated into the proposed logistics architecture and evaluate its effects on survivability, capacity, cost, mobility, and resilience under contested conditions. By advancing from conceptual relay nodes to more capable floating logistics infrastructure, this research agenda aims to refine a practical and scalable model for distributed sustainment in the Western Pacific. For Japan and its allies, the strategic significance of such work lies in building a logistics system that is not only efficient in peacetime, but durable, adaptive, and operationally decisive in crisis and conflict.

References

Blischke, W. R., & Murthy, D. N. P. (2000). Reliability: Modeling, prediction, and optimization. Wiley.

Eccles, H. E. (1959). Operational naval logistics. Bureau of Naval Personnel.

Federal Emergency Management Agency. (2019). National response framework (4th ed.).

Joint Chiefs of Staff. (2019, February 4). Joint publication 4-0: Joint logistics.

Joint Chiefs of Staff. (2022). Joint concept for contested logistics.

Jomini, A.-H. (1862). The art of war (G. H. Mendell & W. P. Craighill, Trans.). J. B. Lippincott. (Original work published 1838)

Ministry of Defense, Japan. (2022a). Defense buildup program.

Ministry of Defense, Japan. (2022b). National defense strategy.

Modarres, M., Kaminskiy, M., & Krivtsov, V. (2016). Reliability engineering and risk analysis: A practical guide (3rd ed.). CRC Press.

Nakagawa, T. (2005). Maintenance theory of reliability. Springer.

Thorpe, G. C. (1986). Pure logistics: The science of war preparation. National Defense University Press.

Toward an AI-Powered Logistics Grid for Maintenance, Repair, and Overhaul: Why the DoW Must Become a Digital OEM to Win the Sustainment War

Dr. Jen Gebhardt—is the Director of Research at Govini, where she leads the integration of data and technology to solve critical defense and industrial base readiness challenges. Previously, as a Product Manager, she spearheaded the development of Govini's defense acquisition software, Ark.ai. Her tenure is defined by deploying scalable, data-driven solutions across the Department of War and federal agencies to mitigate complex national security risks. Formerly an Assistant Professor at Loyola Marymount University, Dr. Gebhardt holds a BA from the University of Maryland and a PhD in Political Science-Political Psychology from the University of California, Irvine.

Abstract

The Department of War (DoW) faces a critical readiness crisis, with billions spent annually on Maintenance, Repair, and Overhaul (MRO) failing to keep advanced weapon systems operational. By examining trends in contract obligations, technical data access, and operational availability across aviation, ground, and maritime systems, this study correlates declining platform availability with low levels of government access to technical data and a shift toward contractor-led MRO. These findings highlight that systemic sustainment challenges are rooted in the government's lack of access to technical data, which precludes competitive sourcing and independent organic repair.

Such reliance is operationally perilous in an era of contested logistics, where distributed sustainment and rapid repair are essential for victory. This paper makes the case that the DoW must recognize sustainment as a data problem and transform MRO through a three-step process: (1) establish the DoW as a Digital OEM through access to technical design files; (2) fuse technical data with AI-enabled software for predictive and prescriptive maintenance; and (3) deploy a decentralized logistics grid leveraging additive manufacturing and distributed sustainment hubs closer to the point of need. Realizing this vision enables the DoW to regenerate combat power at the speed and distance required by modern conflict.

Introduction

The Department of War (DoW) spends billions each year to field and maintain the world's most advanced weapons yet struggles to keep them operational. Aircraft sit grounded, ships remain pierside, and vehicles are cannibalized for parts (U.S. Government Accountability Office [GAO], 2023a, 2025b; U.S. Department of Defense Office of Inspector General [DoDIG], 2023). Not a single tactical aircraft variant met its mission-capable target in Fiscal Year (FY) 2023 (GAO, 2024b). Submarines returning from deployment can wait years for repairs to begin (Eckstein, 2020). Billion-dollar platforms are, too often, idle assets.

This is the state of readiness in peacetime. The situation will compound dramatically in a protracted conflict or contested environment. And the consequences will be much more dire. There are multiple factors that contribute to this crisis and many well-intended attempts to solve it. Yet persistent readiness shortfalls suggest these efforts are treating the symptoms of a deeper, more systemic challenge: Military maintainers and logisticians do not have access to the technical data they need to keep weapons ready and in the fight.

Without access to the technical data, the government cannot competitively source parts, authorize alternative suppliers, or repair systems independently. It outsources not just maintenance, but also control to contractors. The DoW has become a tenant in its own supply chain.

This arrangement is operationally perilous. In an era defined by strategic competition and contested logistics, the difference between victory and defeat will be determined by which side can sustain combat power over time. Modern conflict will not pause for parts to be transported from CONUS depots or for contractor schedules to align with operational requirements. Physical inventory, whether stockpiled or in transit, is itself a vulnerability that will be exploited by our adversaries. The Regional Sustainment Framework was designed with this in mind. Future wars demand distributed maintenance and repair across allied facilities and forward positions (U.S. Department of Defense [DoD], 2024). Yet that vision is impossible when the government lacks the authority to authorize repairs or manufacture spares at the point of need.

The future of readiness is not a warehouse or centralized repair depot. It is not a vulnerable stockpile of physical parts awaiting deployment. It is a secure library of certified digital files, transmissible globally and manufacturable locally. It is smarter data stewardship and utilization. In this model, lead time is measured in data transmission speeds, not shipping schedules. Inventory is theoretically unlimited. The strategic vulnerabilities that undercut the DoW's current maintenance, repair, and overhaul (MRO) model is replaced by the resilience of distributed production.

Realizing this vision requires the DoW to transform its approach to MRO in three decisive steps:

1. **Transform the DoW into a Digital OEM:** Secure access to technical data via digital blueprints as a non-negotiable condition of procurement.
2. **Fuse Data for MRO Intelligence:** Fuse technical data with AI-enabled commercial software to forecast failures, coordinate procurement, and prescribe solutions at machine speed.
3. **Operationalize the Grid:** Deploy a decentralized AI-powered logistics network with additive manufacturing at the point of need—transforming centralized maintenance and repair depots into distributed sustainment hubs.

The Operational Reality an AI-Powered Logistics Grid Enables

A mission-critical hydraulic actuator fails on an F-35 operating from a forward base in the Indo-Pacific. Today, that aircraft is grounded pending retrograde to a CONUS depot or OEM facility—a process measured in weeks or months.

In an AI-powered logistics grid, the response is fundamentally different: Within seconds of the failure, sensor data is compared against a government-owned digital blueprint. An AI model diagnoses the root cause and correlates it against global inventory, supplier capacity, and manufacturing hub availability. The system prescribes the optimal path: additive manufacturing at an allied depot in Japan using a certified digital file. The part is printed, inspected, and installed within 48 hours. The OEM is not consulted. The supply chain is not stressed. The aircraft returns to mission-capable status at a fraction of the traditional cost and timeline.

The future of MRO lies not in more funding or better wrench-turning, but in smarter data ownership and utilization. For decades, the DoW has wrestled with sustainment challenges while the underlying incentives and data structures remained unchanged. The DoW has a choice. It can continue managing sustainment as a logistics problem, or it can recognize sustainment as a data problem and seize the authority required to solve it.

The cost of inaction is both financial and strategic. Sustainment consumes roughly 70% of a weapon system's total life-cycle cost—meaning the DoW spends twice as much to maintain these platforms as it did to procure them (GAO, 2024a). Yet expensive platforms that cannot be sustained cannot deter and cannot be used by the Joint Force. In a future conflict that will be defined by speed and distance, the side that can repair and regenerate combat power at the point of need holds a strategic advantage.

The Modern Sustainment Challenge

The United States possesses the world's most robust Organic Industrial Base (OIB)—a network of depots and arsenals staffed by skilled maintainers who form the backbone of warfighting readiness. Beyond the OIB, thousands of maintainers and sustainers keep the force operational day-to-day. Yet neither workforce is equipped with the data and data infrastructure they need to effectively do their jobs. Consequently, the department has been forced to rely on an expensive, contract-heavy MRO model.

Since FY2016, the DoW has spent over $286 billion for MRO *in contract obligations alone*—an average of $29.4 billion each year.[1] Importantly, these figures exclude costs associated with the bulk of the DoW's Operations and Maintenance budget, such as government personnel, facility construction and maintenance, base support, training, and day-to-day operations. To put these figures in perspective, the global commercial aviation sector, with a fleet size of around 30,000 aircraft, annually averaged $80 billion on MRO over this time period (Oliver Wyman, 2022, 2024). This means the DoW's spending on MRO awards alone amounts to over a third of the size of the entire commercial aviation MRO market. Still, this enormous investment has not bought ready systems and platforms.

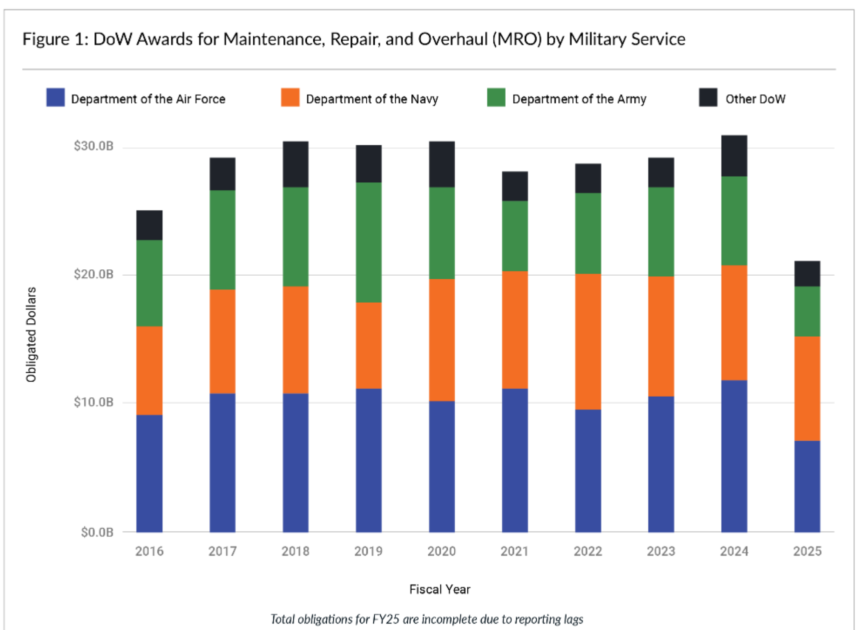

Figure 1. DoW Awards for MRO by Military Service

[1] Numbers represent obligated dollars via contract awards and, therefore, are not inclusive of all DoW MRO spend, such as costs associated with governmet personnel, construction and facility maintenance, transportation, capital equipment, and other depot operating costs.

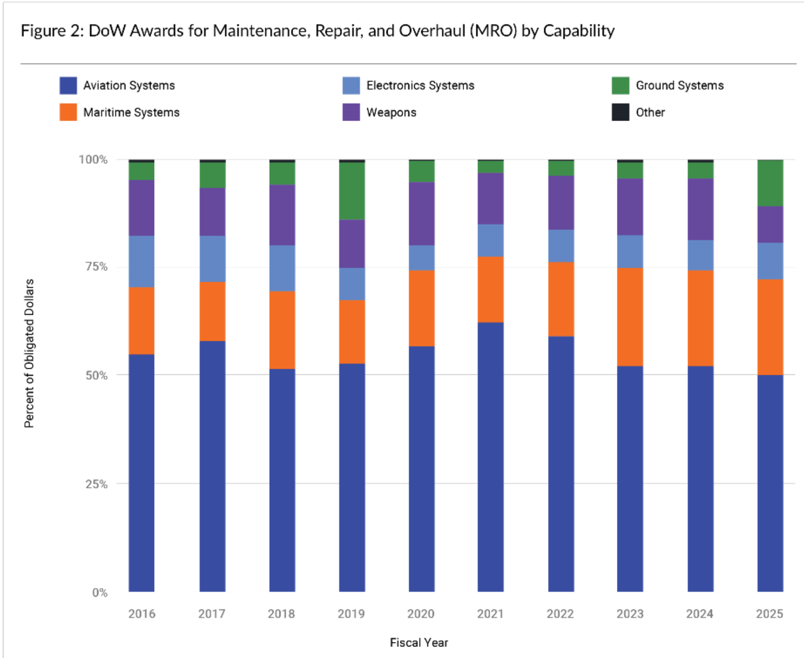

Figure 2. DoW Awards for MRO by Capability

Instead, mission-capable rates have declined, particularly among aviation systems, which account for more than half of all MRO dollars awarded each year. Half of the aircraft reviewed by the GAO (2022a, 2023a) failed to meet their annual mission-capable goal in any fiscal year between FY2011 and FY2021. For tactical aircraft, the situation is particularly critical: None of the aircraft variants reviewed met their mission-capable goals in FY2023 (GAO, 2024b). Critical fifth-generation platforms, such as the F-35, have struggled significantly, failing to meet reliability and maintainability goals and resulting in operational availability rates around 55% against a 90% objective (GAO, 2023d).

In peacetime, this is a budgetary inefficiency. In wartime, it is a critical failure point.

Increased Reliance on Contractors and Decreased Reliance on the Organic Industrial Base for MRO

In recent years, the DoW's response to sustainment complexity has been to delegate it via Contractor Logistics Support (CLS), Performance-Based Logistics (PBL), and similar life-cycle sustainment contracts (Light et al., 2024). In doing so, the DoW seeks to transfer risk and leverage specialized industry expertise to maintain its platforms and weapons systems. While intended to improve readiness, these contracts deepen the DoW's dependency on OEMs and obscure visibility into the sub-tier supply chain. They further reduce government access to the underlying technical data, making it difficult to validate costs, manage risks, or introduce competition. The result is that the DoW often outsources sustainment while retaining the ultimate risk of failure. This structural reliance prevents the organic base from developing the institutional muscle memory required to repair systems independently during a crisis.

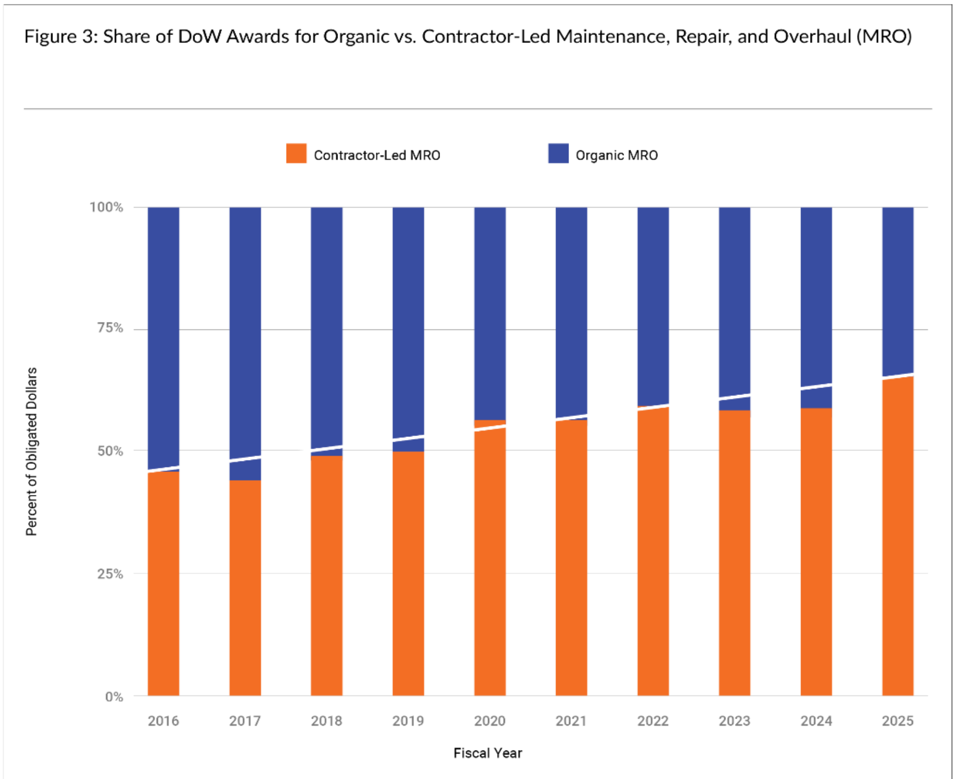

Figure 3. Share of DoW Awards for Organic vs. Contractor-Led MRO

Figure 3 reveals a decisive shift toward contractor-led sustainment, with contracted MRO (including awards for CLS and PBL) growing at a Compound Annual Growth Rate (CAGR) of 5.7%. Dollars awarded for organic MRO stagnated over this time but declined relative to all MRO dollars awarded. In FY2016, less than half of all dollars went to contractor-led MRO. In FY2025, 64% of awarded dollars went to contracted MRO. This trend reveals a self-reinforcing cycle: As contractor dependency grows, the government further loses access to the data, systems, and supply chain visibility it needs to sustain its own weapons.

Data Fragmentation and Systemic Silos

The foundation of any predictive enterprise is unified, high-quality data. However, the DoW's MRO ecosystem is built on a patchwork of legacy systems that are inoperable across programs, platforms, and services. Critical data—from part failure histories and supplier lead times to depot capacity and mission capable rates—reside in disconnected silos (GAO, 2022b). This fragmentation prevents the aggregation and correlation necessary for accurate readiness forecasting.

The F-35's Autonomic Logistics Information System (ALIS) exemplifies this challenge. Since FY2010, the DoW has awarded Lockheed Martin $2.4 billion to sustain and modernize ALIS. Despite this investment, the system has been publicly criticized for being frustrating and inefficient, requiring maintainers to use manual workarounds for tasks that should be automated (Clark, 2019). Data fidelity and usability issues became so severe that ALIS required a complete re-architecture into the Operational Data Integrated Network (ODIN) at additional cost (Congressional Research Service [CRS], 2024).

This is not an isolated case; it is emblematic of a defense-wide struggle to establish a standardized, unified data foundation to optimize and use data at scale. The DoW's perennial struggle to achieve sound financial management is similarly tied to the fragmentation and poor oversight of the underlying business systems that manage assets and logistics (GAO, 2023b). In this environment, even the best-intentioned CLS or PBL arrangement struggles to achieve true efficiency because it lacks the necessary data for predictive forecasting.

Technical Data Authority: The Missing Foundation for Competitive Sustainment

Data fragmentation and increased reliance on contractor-led MRO are symptoms of a deeper structural challenge: The government lacks access to the technical data underlying its weapon systems. Contractors legitimately protect their intellectual property, but without access to technical data, the government cannot authorize alternative suppliers, validate costs, or repair independently (GAO, 2025a). This dependency creates a structural barrier to competitive sustainment and operational flexibility.

Often, when a critical component fails, DoW maintainers are not authorized to fabricate a replacement or conduct the repair themselves without the technical data (DoDIG, 2022). They must wait for a vendor to respond or resort to cannibalizing other platforms, a practice that increases maintenance workload, degrades the long-term health of the fleet, and creates significant risks for long-term operational availability (GAO, 2023d). It also means that the government often lacks alternative sources to fill supply gaps.

Air Force Secretary Frank Kendall cited the F-35 as a "painful" lesson in this regard, noting that its "total system procurement" model ceded intellectual property rights to Lockheed Martin, effectively preventing the government from cost-effectively upgrading the platform or leveraging competition for repairs (*Posture of the Department of the Air Force*, 2024).

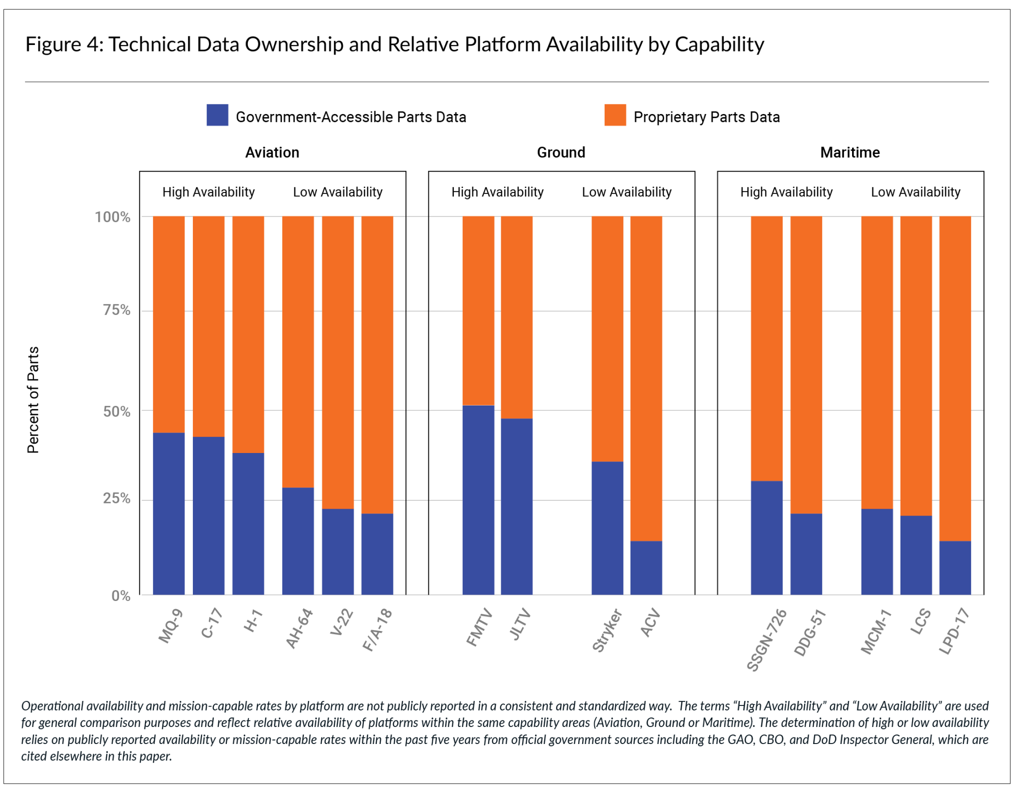

Figure 4. Technical Data Ownership and Relative Platform Availability by Capability

Figure 4 reveals the pervasive lack of government access to technical data across Aviation, Ground, and Maritime capabilities. Across all platforms in Figure 4—whether high or low availability—the government lacks access to technical information for the majority of parts. In most cases, government-accessible part data is well below 50%. In addition, programs with lower availability tend to correlate with a lower degree of government data access. The F/A-18 Super Hornet, Stryker Combat Vehicle, and Littoral Combat Ship (LCS) exhibit some of the lowest levels of government data access. Such platforms also have particularly low operational availability or mission-capable rates (CBO, 2022; GAO, 2022a, 2023a, 2024b, 2025b).

MRO on the F/A-18 Super Hornet is a clear illustration of this problem. The Navy could not repair or establish an alternative source for generator converter units because the contractor would not share the necessary technical data and would not respond to a request for pricing the data package (DoDIG, 2019). Similarly, the DLA could not acquire a specific communication antenna because the government could not access technical drawings, and the sole-source manufacturer was experiencing a 13-month production delay, making the part impossible to source (DoDIG, 2019).

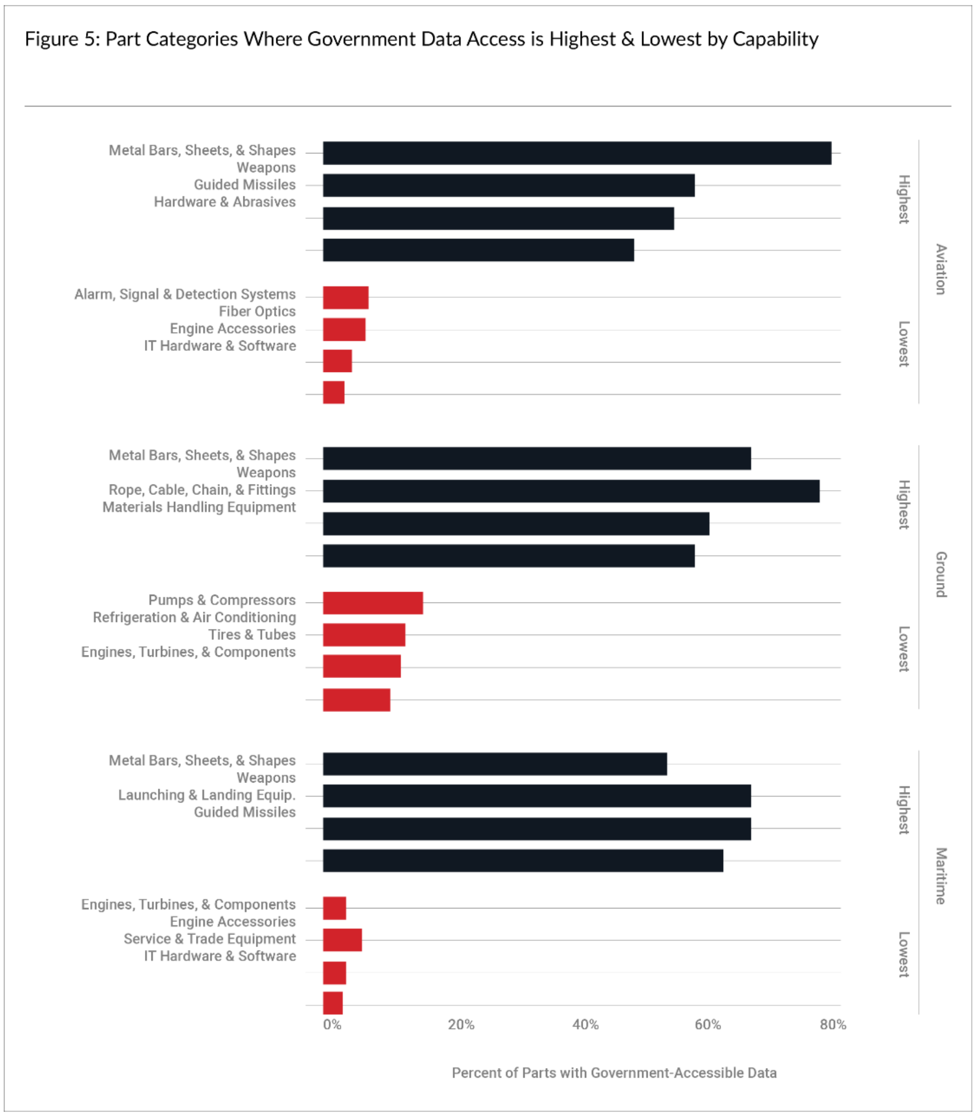

Figure 5. Part Categories Where Government Data Access is Highest & Lowest by Capability

The problem is most acute where it matters most. As shown in Figure 5, the government's access to data is lowest for the most complex and technologically advanced components. While the DoW may have access to data for basic structural materials, it consistently lacks access to critical subsystems such as engines and avionics. When the government attempts to obtain access to this proprietary data, which is often owned by subcontractors rather than the prime contractor, it can be met with refusal or extremely high prices for the data rights (Light et al., 2024). This is the technical equivalent of owning the blueprint for a car's frame, but not for its engine or transmission, making independent repair or competitive sourcing for the most vital parts impossible. As a result, parts become backordered, and platforms sit idle.

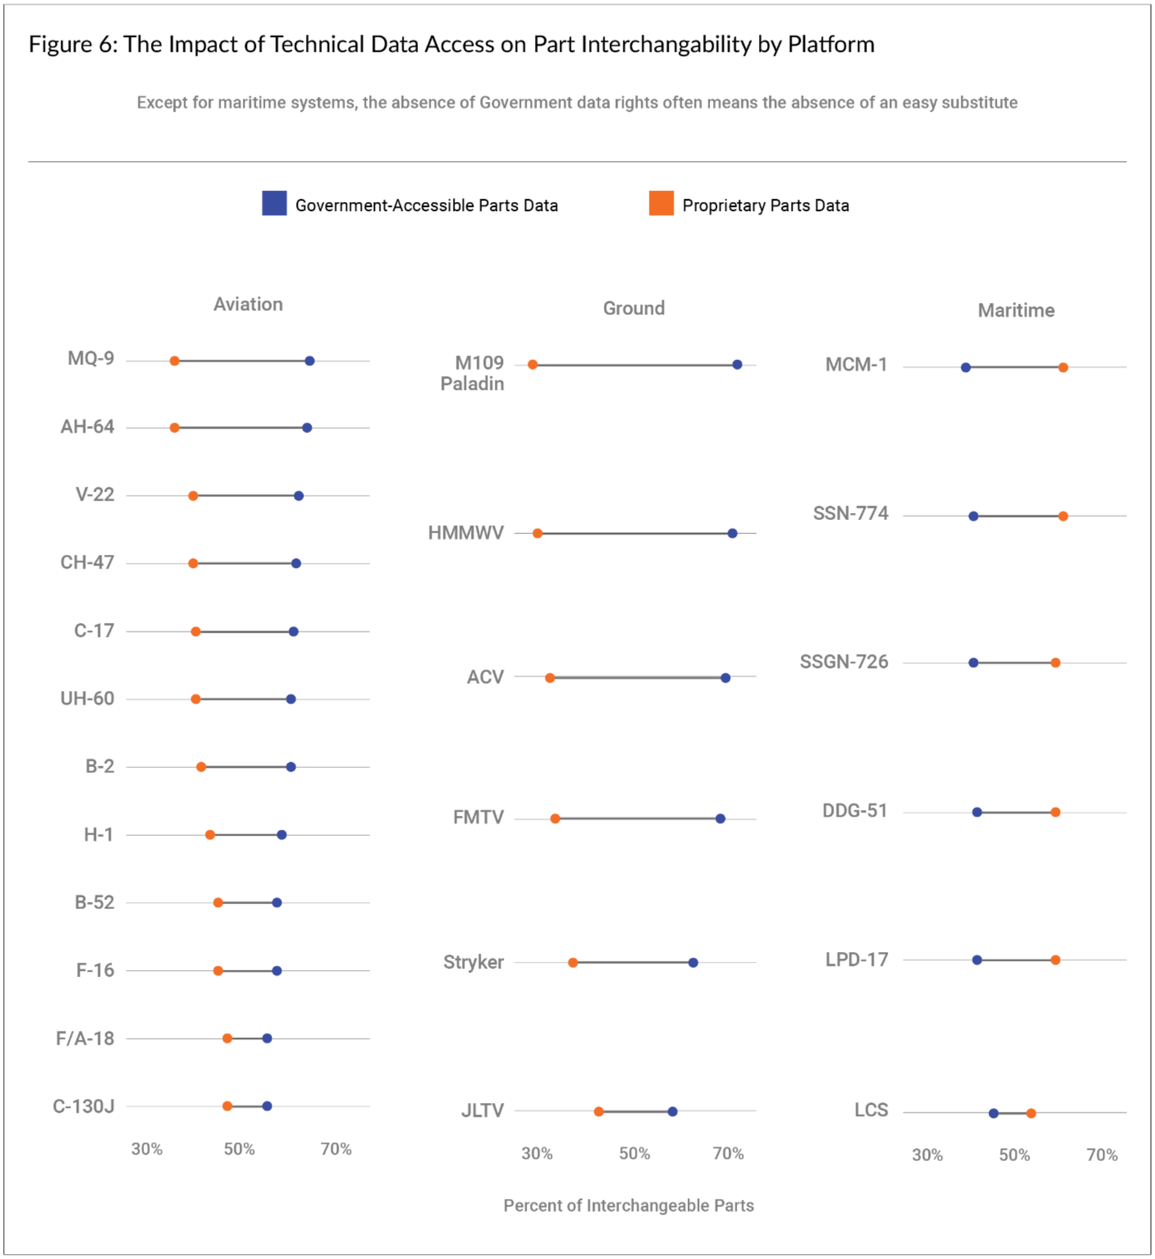

Figure 6. The Impact of Technical Data Access on Part Interchangeability by Platform

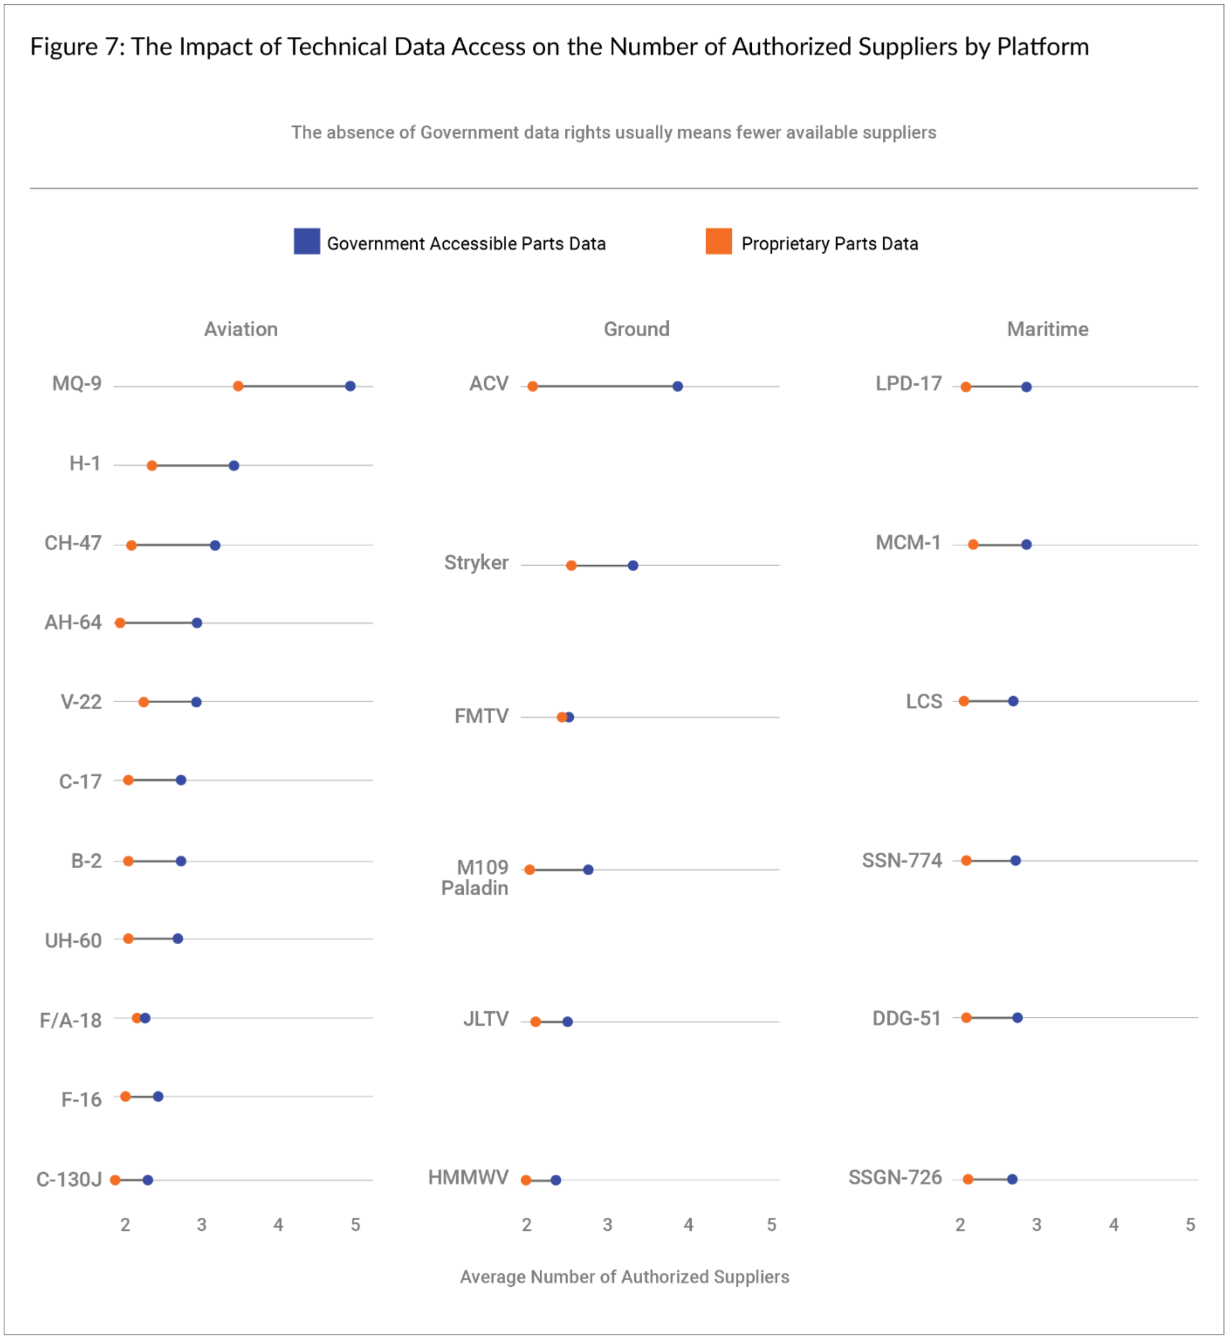

Figure 7. The Impact of Technical Data Access on the Number of Authorized Suppliers by Platform

When the government lacks access to technical data, it often lacks alternative parts and alternative suppliers to fill supply gaps. Figure 6 shows that when the government has access to a part's data, there is a much greater chance that an interchangeable substitute is available. Conversely, when a part's data is proprietary and inaccessible to the government, it is less likely that it can be swapped with an interchangeable substitute. The exception to this finding is for Maritime systems, which suggests that even with government-accessible data, the unique and highly specialized nature of naval components makes finding substitutes difficult.

Figure 7 reveals a consistent pattern across Aviation, Ground, and Maritime platforms: Parts with government-accessible data have a greater number of authorized suppliers than their proprietary counterparts. The reverse is also true. This lack of competition is not a failure of logistics; it is a failure of acquisition strategy that reverberates for decades.

Transforming the MRO Enterprise

The DoW must address the cause of persistent sustainment problems, not just the symptoms. The DoW cannot fix MRO one depot or one contract at a time. Incremental reforms have failed because they address symptoms rather than the underlying structure. True transformation requires three integrated steps that fundamentally reorder the relationship between the government, its data, and the industrial base.

Step 1: Transform the DoW into a Digital OEM

Transforming the MRO enterprise begins with data. The DoW's 2025 Acquisition Transformation Strategy recognizes that the government cannot deliver capabilities to the American warfighter at pace to deter and defeat our adversaries without greater access to and utilization of the technical data underlying its systems. It calls for "owning the operator's manual of our systems so government and military mechanics can repair our weapon systems without reliance on others," which ensures that organic depots and contractors alike can repair systems (DoW, 2025). Members of Congress have attempted to address this challenge with "right to repair" amendments to the NDAA with limited success (GAO, 2025a; Hultz, 2025). A more balanced and sustainable approach is needed that will empower the DoW to improve the speed and flexibility by which it can deliver for the warfighter.

Securing access to technical data via digital blueprints accomplishes three critical goals:

1. **It Breaks Vendor Lock:** It enables robust competition for sustainment, allowing organic depots and innovative commercial partners to repair, replace, and upgrade components.
2. **It Empowers the Organic Industrial Base**: It gives government depots the information they need to fulfill their statutory mandate, turning them into centers of excellence and innovation.
3. **It Creates the Foundation for an AI-Powered Logistics Grid:** It provides the standardized, complete, and trusted data that is the absolute prerequisite for systematic predictive maintenance or AI-driven logistics.

Step 2: Fuse Data for MRO Intelligence

Once the DoW has access to the core technical data, AI-enabled commercial software must fuse it with real-world procurement, supplier, and operational intelligence to unlock predictive and even prescriptive capabilities. MRO intelligence is the output of this fusion, unlocking a range of capabilities, including:

- **Predictive Maintenance Planning:** AI models can instantly forecast maintenance demand from usage data and prescribe the optimal sustainment path, eliminating human lag time. An AI model could have prospectively identified, for example, the long-term consequences of the F-35's engine cooling problem. The Pentagon knew the F135 engine was operating beyond its design parameters but accepted an estimated $38 billion in future maintenance costs to avoid initial program delays (GAO, 2023c). An AI-powered analysis could have quantified that future risk in today's dollars, arming program managers with the evidence to demand a design change.

- **Cross-Program Optimization:** AI can identify shared components, suppliers, and risks across hundreds of programs—turning fragmented maintenance into system-level insight. This system-level insight allows the DoW to manage its industrial base as a strategic portfolio, anticipating bottlenecks and coordinating procurement in ways that are impossible today.

This integrated intelligence transforms maintenance from a slow, reactive sustainment function into a readiness driver, enabling true predictive readiness and supporting distributed operations.

Step 3: Operationalize an AI-Powered Logistics Grid for Contested Operations

By establishing authority as a Digital OEM and integrating MRO intelligence, the DoW can build the sustainment enterprise required for modern conflict: a decentralized logistics grid and MRO network. Leveraging MRO intelligence, the grid can autonomously manage a global network of certified industrial bases and additive manufacturing hubs. It can determine whether a part should be repaired at a U.S. or allied depot, procured from pre-positioned stock, or additively manufactured at the point of need using a government-accessible digital data file.

This is not speculative. The technology exists (Charkaluk & Szmytka, 2025). What has been missing is the data access to deploy it. By becoming a Digital OEM, the DoW unlocks this capability across the entire sustainment enterprise.

From Physical Inventory to Digital Readiness

Becoming a Digital OEM means asserting, through policy and acquisition practice, that the government must have access to the complete technical data package for every weapon system it procures. This does not suggest that the government must mandate this access for free, but it means refusing to pay twice—once for the weapon and then later for the privilege of maintaining it. Realizing this vision does not require a wholesale upheaval of the defense industrial base's business models or a complete dismissal of legitimate intellectual property rights. Instead, existing frameworks—specifically within Title 10—present a logical starting point for modernizing the DoW's approach to MRO. Updating how the law treats technical data authority within these statutes is a pragmatic first step toward the Digital OEM model, ensuring that the transition to a distributed logistics grid is both strategically decisive and respectful of the vital partnerships that define the U.S. industrial base.

Above all, any legislative or policy changes must reflect the 21st century reality that access to critical data is as strategically vital as physical infrastructure or exquisite lethality.

With this foundation of data, AI-enabled software becomes the engine of transformation. By fusing technical blueprints with maintenance data and real-world supply chain intelligence, the DoW can finally shift from reactive repair to predictive readiness. It can identify risks before they cascade into fleet-wide groundings. It can dynamically route maintenance to the optimal location—whether that's an organic depot, an allied facility, or an additive manufacturing hub embedded with forward forces. It can make decisions at machine speed that currently require weeks of coordination.

This is the vision that makes the Regional Sustainment Framework operationally viable. In a protracted conflict, the victor will not be the side with the largest stockpiles, but the side that can regenerate combat power faster than it is destroyed. An AI-powered logistics grid, underpinned by government-accessible technical data and best-in-class commercial software, turns that requirement from aspiration into reality.

References

Charkaluk, É., & Szmytka, F. (2022, May 31). *3D printing is set to hit the battlefield*. Polytechnique Insights. https://www.polytechnique-insights.com/en/columns/science/3d-printing-is-set-to-hit-the-battlefield/

Clark, C. (2019, March 6). *Air Force tries to fix F-35's ALIS — From a big, broken box to the cloud*. Breaking Defense. https://breakingdefense.com/2019/03/air-force-moving-f-35s-alis-from-a-big-broken-box-to-the-cloud/

Congressional Budget Office. (2022, January). *Availability and use of aircraft in the Air Force and Navy*. https://www.cbo.gov/publication/57433

Congressional Research Service. (2024, December 3). *F-35 Lightning II: Background and issues for Congress* (CRS Report No. R48304). https://crsreports.congress.gov/product/pdf/R/R48304

Eckstein, M. (2020, May 26). *NAVSEA says attack sub repairs much improved as USS Boise enters yard following 4-year wait*. USNI News. https://news.usni.org/2020/05/26/navsea-says-attack-sub-repairs-much-improved-as-uss-boise-enters-yard-following-4-year-wait

Hultz, C. (2025, October 28). *Military right to repair reform could reshape readiness and innovation*. Military.com. https://www.military.com/feature/2025/10/24/military-right-repair-reform-could-reshape-readiness-and-innovation.html

Light, T., Chien, C., & Sanchez, R. (2024). *Management of U.S. Air Force aircraft contractor logistics support arrangements* (Report No. RR-A194-1). RAND. https://www.rand.org/pubs/research_reports/RRA194-1.html

Oliver Wyman. (2022, February). *Executive summary: 2022 global fleet & MRO market forecast*. https://www.oliverwyman.com/our-expertise/insights/2022/feb/global-fleet-and-mro-market-forecast-2022-2032.html

Oliver Wyman. (2024). *Global fleet and MRO market forecast 2024-2034*. https://www.oliverwyman.com/our-expertise/insights/2024/feb/global-fleet-and-mro-market-forecast-2024-2034.html

Posture of the Department of the Air Force in review of the Defense Authorization request for Fiscal Year 2025 and the future years: Hearing before the Committee on Armed Services, 118th Cong. (2024). https://www.congress.gov/event/118th-congress/senate-event/335690

U.S. Department of Defense. (2024, May 9). *Regional Sustainment Framework (RSF)*. https://www.acq.osd.mil/asds/docs/RSF-9MAY24.pdf

U.S. Department of Defense Office of Inspector General. (2019, November 19). *Audit of Navy and Defense Logistics Agency spare parts for F/A-18 E/F Super Hornets* (Report No. DODIG-2020-030). https://www.dodig.mil/reports.html/Article/2022428/audit-of-navy-and-defense-logistics-agency-spare-parts-for-fa18-ef-super-hornet/

U.S. Department of Defense Office of Inspector General. (2022, June 21). *Audit of the Department of Defense's implementation of predictive maintenance strategies to support weapon system sustainment* (Report No. DODIG-2022-103). https://www.dodig.mil/reports.html/Article/3063635/audit-of-the-department-of-defenses-implementation-of-predictive-maintenance-st/

U.S. Department of Defense Office of Inspector General. (2023, May 23). *Management advisory: Maintenance concerns for the Army's prepositioned stock–5 equipment designated for Ukraine* (Report No. DODIG-2023-076). https://www.dodig.mil/reports.html/Article/3407150/management-advisory-maintenance-concerns-for-the-armys-prepositioned-stock5-equ/

U.S. Department of War. (2025, November). *Acquisition transformation strategy*. https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/ACQUISITION-TRANSFORMATION-STRATEGY.PDF

U.S. Government Accountability Office. (2022a, November 1). *Weapon system sustainment: Aircraft mission capable goals were generally not met and sustainment costs varied by aircraft* (Report No. GAO-23-106217). https://www.gao.gov/products/gao-23-106217

U.S. Government Accountability Office. (2022b, December 8). *Military readiness: Actions needed to further implement predictive maintenance on weapon systems* (Report No. GAO-23-105556). https://www.gao.gov/products/gao-23-105556

U.S. Government Accountability Office. (2023a, January 31). *Weapon system sustainment: Navy ship usage has decreased as challenges and costs have increased* (Report No. GAO-23-106440). https://www.gao.gov/products/gao-23-106440

U.S. Government Accountability Office. (2023b, March 7). *Financial management: DOD needs to improve system oversight* (Report No. GAO-23-104539). https://www.gao.gov/products/gao-23-104539

U.S. Government Accountability Office. (2023c, May 30). *F-35 Joint Strike Fighter: More actions needed to explain cost growth and support engine modernization decision* (Report No. GAO-23-106047). https://www.gao.gov/products/gao-23-106047

U.S. Government Accountability Office. (2023d, September 21). *F-35 aircraft: DOD and the military services need to reassess the future sustainment strategy* (Report No. GAO-23-105341). https://www.gao.gov/products/gao-23-105341

U.S. Government Accountability Office. (2024a, February 29). *Weapon system sustainment: DOD identified operating and support cost growth but needs to improve the consistency and completeness of information to Congress* (Report No. GAO-24-107378). https://www.gao.gov/products/gao-24-107378

U.S. Government Accountability Office. (2024b, October 21). *Tactical aircraft: Operation and maintenance spending varies by system, and availability generally does not meet service goals* (Report No. GAO-25-107870). https://www.gao.gov/products/gao-25-107870

U.S. Government Accountability Office. (2025a, September 29). *Weapon system sustainment: DOD can improve planning and management of data rights* (Report No. GAO-25-107468). https://www.gao.gov/products/gao-25-107468

U.S. Government Accountability Office. (2025b, September 2). *Weapon system sustainment: Various challenges affect ground vehicles' availability for missions* (Report No. GAO-25-108679). https://www.gao.gov/products/gao-25-108679

From Factory to Field: Modeling Production Capacity and Logistics Effectiveness for Defense System

Nikita Birbasov-Drosendahl—is a doctoral candidate at the Georgia Institute of Technology with a master's degree in aerospace engineering from Georgia Institute of Technology and a bachelor's degree in mechanical engineering from University of Kentucky. [nbirbasov3@gatech.edu]

Paul Wang—is a doctoral candidate at the Georgia Institute of Technology. His research is on data-driven modeling methods for aerospace systems design. [pwang395@gatech.edu]

Dr. Olivia Pinon Fischer—is a Principal Research Engineer at 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@ae.gatech.edu]

Dr. Michael Belisle—is Staff Product Architect, Advanced Development, at Northrop Grumman Corporation. [michael.belisle@ngc.com]

Dr. Dimitri Mavris—is a Georgia Tech Distinguished Regents Professor, Director of the Aerospace Systems Design Laboratory (ASDL), and Executive Director of the Professional Master's in Applied Systems Engineering (PMASE). He is a Fellow of the American Institute of Aeronautics and Astronautics (AIAA) and Royal Aeronautical Society. [Dimitri.mavris@ae.gatech.edu]

Abstract

Recent conflicts have underscored the importance of adaptable supply lines in contested logistics and rapidly evolving technological environments. Traditional supply chains respond slowly to changing conditions and often lack traceability to campaign-level measures of performance. To address this gap, we propose a modeling and simulation (M&S) environment that supports data-driven assessment of alternative scenarios. A previous phase of the effort developed an overall logistics model for the delivery and storage of goods, although that model has not yet been fully integrated with the present work. This paper focuses on a higher-fidelity model of a notional manufacturing facility and evaluates how production-line design changes affect campaign-level supply-chain performance. The manufacturing steps are based on publicly available descriptions of Ling-Temco-Vought's Multiple Launch Rocket System production process. Two scenarios are examined: a battle scenario that transitions from high-demand engagement to lower-demand sustainment and a second scenario that introduces an unexpected surge in demand.

Introduction

Military logistics has emerged as a critical capability for the United States as it prepares to compete with near-peer adversaries. Although frontline platforms often receive the most attention and investment, sustaining those forces over time remains a persistent challenge. Recent efforts to supply Ukraine have highlighted vulnerabilities in the Defense Industrial Base (DIB) on which U.S. forces would depend in a large-scale conflict. Supporting a modern fighting force requires long production lead times, early contractual commitments, and clear visibility into demand across the entire supply chain, from the factory floor to the point of use in theater. More broadly, the current era of Great Power Competition is forcing a reassessment of DIB readiness. As noted in the NDIA & DRIVE Contested Logistics Workshops Summary & Recommendations, "The readiness of the U.S. defense enterprise to sustain potential & prolonged conflicts is increasingly a concern" (National Defense Industrial Association [NDIA], 2025). In this context, the DIB must evolve into a threat-informed defense ecosystem that can scale production, satisfy surge demand, and recover during a major conflict. Among the resulting recommendations, improvement in manufacturing planning and execution stands out as a clear priority.

This research investigates how manufacturing capacity for a missile launcher affects the ability to satisfy demand across different conflict intensities and tempos. The effort has three objectives: (1) develop a digital modeling environment for launcher manufacturability and production-capacity analysis, (2) conduct sensitivity analyses to determine how process parameters, including machine capacity, affect throughput, and (3) quantify how alternative factory configurations influence the ability to meet specific demand signals.

The Background and Methodology sections first describe the broader factory-to-field modeling effort and the architecture that links demand, logistics, and production. The remainder of the paper focuses on the manufacturing-facility model developed in this phase, the surrogate-modeling approach used to analyze it, and the resulting implications for production-capacity planning.

Background

The overall effort couples logistics and production-system models to create a factory-to-field simulation environment. At the highest level, the logistics model is intended to evaluate supply and transportation capacity, determine how those factors affect mission outcomes, and identify the variables that most improve mission completion. Discrete-event simulation has been widely used for this purpose. Cantrell et al. (2025) used a discrete-event simulation to assess the impact of digital twins under conditions of part degradation, changing conflict intensity, and varying resupply or repair availability. Hoem et al. (2025) examined the effect of predictive-maintenance technologies enabled by digital twins by building a discrete-event simulation of an aircraft logistics system. Metcalf and Laffey (2023) proposed a digital-twin–enabled decision-support system for wargaming and demonstrated the value of such models for what-if analysis when assessing technology or strategy infusions. Paksoy et al. (2012) developed a nonlinear mixed-integer model that integrates supply-chain design and factory line balancing; however, their formulation assumes deterministic parameters and constant demand throughout the scenario.

The overall logistics system is composed of multiple interacting facilities. Because the long-term goal is to simulate the production-to-field process, appropriate factory representation and evaluation methods are also needed. This aspect is directly informed by the work presented in Sieldak et al. (2018), Pinon Fischer et al. (2022), Siedlak et al. (2015), and Libby et al. (2017).

Methodology

This section presents the methodology used to evaluate how manufacturing capacity influences the ability of the broader supply chain to satisfy operational demand. The overall approach links demand, logistics, and production models within a common simulation environment, while the present work focuses specifically on the manufacturing facility component. To support this analysis, manufacturing process times are estimated in SEER-MFG® and then incorporated into a SIMIO® discrete-event model to assess throughput, bottlenecks, and the effects of alternative factory configurations under different demand conditions.

Proposed Architecture of the Overall Environment

To gain full visibility into demand across the supply chain, three integrated modeling capabilities are required, as shown in Figure 1: a demand simulation (or battle simulation) to forecast material drawdowns and delivery locations, a logistics simulation to assess transportation capacity and warehouse status across the theater, and a manufacturing or production model to estimate throughput. Integrating these models enables end-to-end analysis of the supply chain and supports the adjustments needed to meet operational demand. The overall approach is shown in Figure 2.

The logistics simulation environment was established in prior work (Birbasov et al., 2025). The present effort focuses on developing the manufacturing-process and production-facility models needed to estimate throughput for alternative factory configurations. Specifically, the study models the steps involved in manufacturing the Ling-Temco-Vought (LTV) Multiple Launch Rocket System (MLRS®) using publicly available information on the production process (LTV Aerospace and Defense, n.d.). Processing times for the relevant manufacturing activities are estimated in SEER-MFG, while the factory floor, including the type and number of machines, is modeled in Simulation Modeling framework based on Intelligent Objects (SIMIO), a discrete-event simulation environment, to estimate resulting throughput.

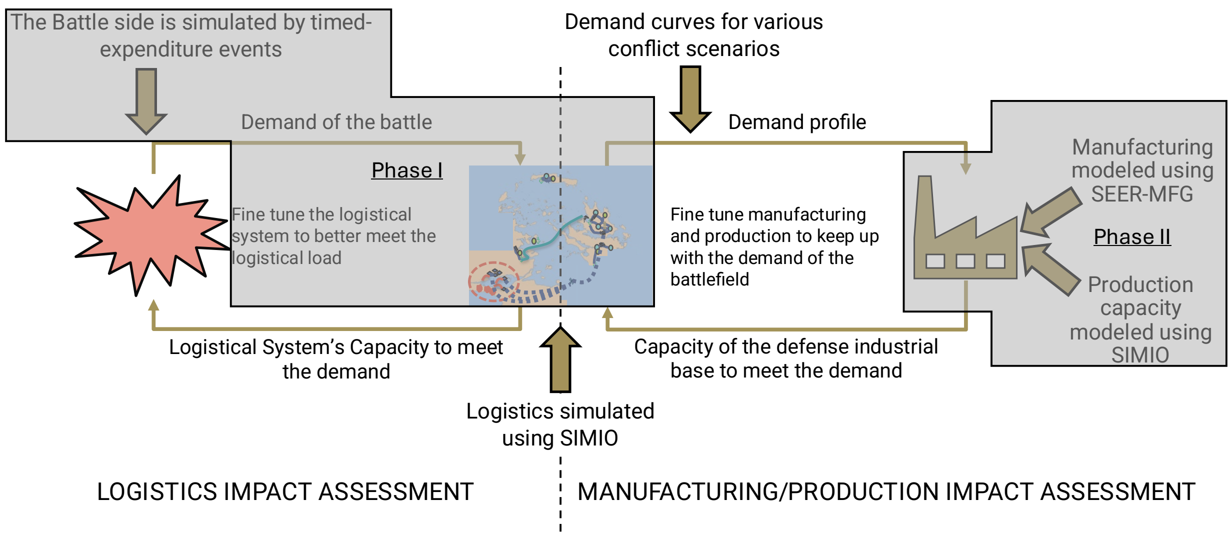

Figure 2. The Overall Proposed Environment. The grey boxes show work done during Phase I and Phase II of the implementation.

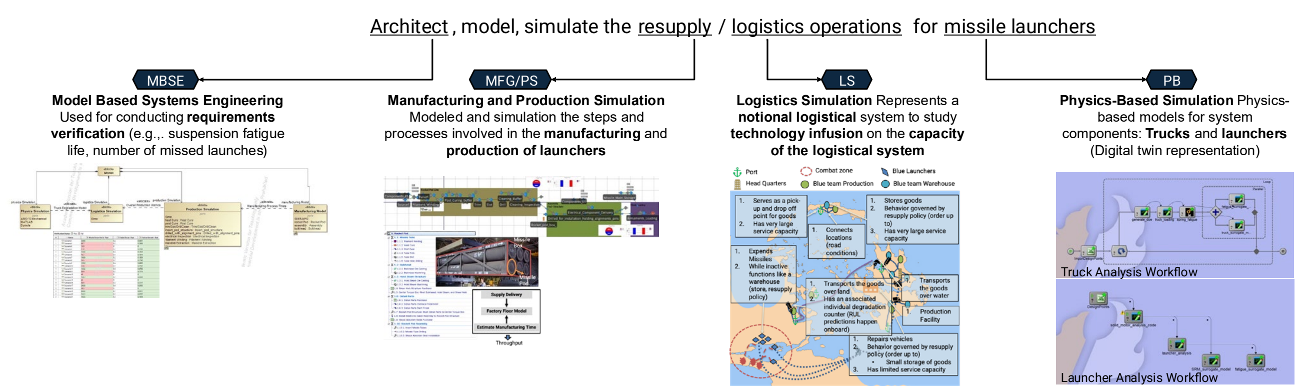

Figure 3. The Overall Simulation Approach Steps

The following sections detail the development of the manufacturing model and the creation of the production model.

Manufacturing Model

This section explains how process-level manufacturing estimates are generated and then translated into a factory-level simulation. Together, the SEER-MFG and SIMIO models provide the analytical link between part design, facility configuration, and production throughput.

Manufacturing Process Modeling

Manufacturing processes are modeled in SEER-MFG (Galorath, n.d.) to estimate the time required for each activity on the production line. Each manufacturing activity is represented in SEER-MFG with its own set of parameters, such as material, geometry, hole count, and machining details. Users can modify these inputs and run the model to obtain process-time estimates derived from a historical database of manufacturer data. In this work, the complete production line is represented in SEER-MFG.

Not all parameters are specified directly by the user. Any inputs that are not explicitly modified are populated by the software's knowledge bases, which are derived from a database of commercial manufacturing processes. These knowledge bases function as process templates that assign values to the remaining parameters within each activity. For example, the composite-manufacturing module includes a filament-winding knowledge base, which is used here to define the filament-winding step.

This work evaluates how changes in selected design variables affect process times and, in turn, downstream logistics performance. In particular, the study varies missile-tube geometry and missile-tube material, then runs SEER-MFG to estimate the process time for each geometry-material combination. The full set of input parameters is listed in Table 1.

Table 1. SEER-MFG Input Parameters

Parameter	Input(s)
Missile Tube Geometry	Diameter [in]
	Length [in]
	Thickness [in]
Missile Tube Material	Fiber Filament
	Weight [lb]

These parameters are used to edit a set of input files for each manufacturing process. Templates for these files are obtained by exporting commands to run the model from SEER-MFG as text files, which are edited with user parameters to create the input files.

SEER-MFG is executed in server mode to automate batch runs. A wrapper script launches the software and continuously queries SEER-MFG for process-time results over a set of input values. This automation supports a design of experiments that captures how process times vary with tube geometry, material, and weight. The procedure is summarized in Figure 4.

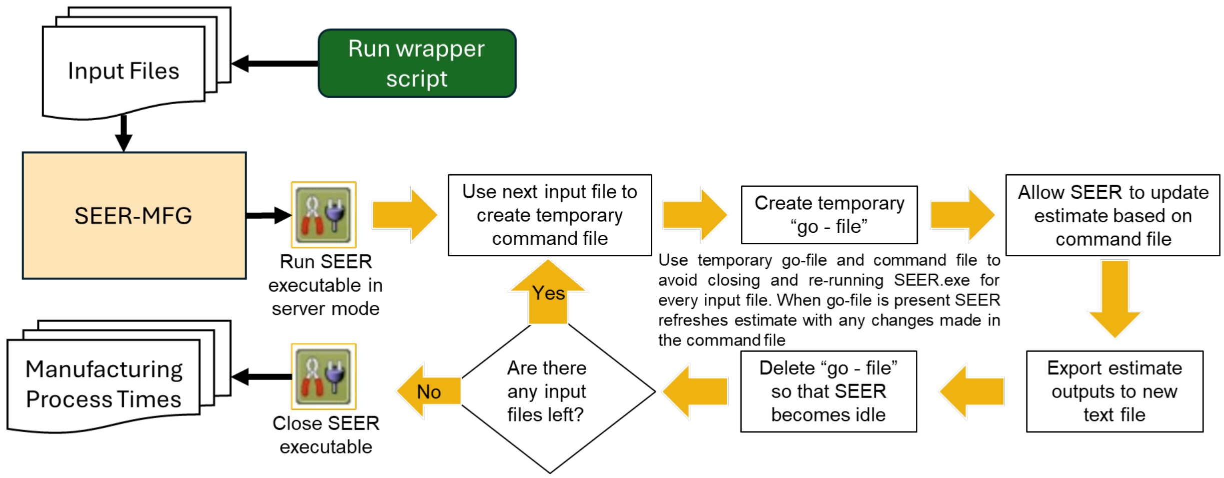

Figure 4. Process of Running SEER-MFG Through Server Mode

Once all runs are complete, the resulting manufacturing-process times can be passed to the factory model to assess their effects on overall production throughput.

Factory Modeling

The purpose of the factory (production) model is to represent a notional manufacturing system in which users can test how technology infusions and factory-floor adjustments affect the industrial base's ability to meet demand. As noted above, the factory model is implemented in SIMIO. SIMIO is an object-based modeling platform that can be extended to multiple paradigms, including agent-based, discrete-event, and continuous simulation. In SIMIO, each object is defined by processes, events, and tables that govern its behavior. The following subsections describe the model elements, along with the logic used to represent production steps and supply deliveries.

The first general assumption is that manufacturing tools and machines do not degrade over time; accordingly, production stoppages due to equipment failures are not modeled.

The second general assumption is that sufficient labor is always available for the planned production shifts.

The third assumption is that the factory operates continuously, 24 hours per day and 7 days per week, without interruptions due to shift changes.

Each subsection below follows a similar format: discussion of the desired modeling task, assumptions, and implementation logic.

The Production Floor

The production floor was based on publicly available information about LTV's MLRS manufacturing facility. Because detailed data on building size and the space required for each production step were not available, a notional facility was constructed using the manufacturing steps identified for LTV's production floor. The resulting layout includes a rocket-pod production line, an assembly line, and an armament-loading line, as summarized in Table 2.

Table 2. Manufacturing Steps

Rocket Pod Line	Assembly Line	Armament Loading Line
1. Filament Winding	1. Insertion of Pods into a Structure	1. Loading Missiles into the Assembled Pod Structures
2. Heat Curing	2. Installation of Holding Alignment Pins	
3. Forming Mandrel Extraction	3. Electrical Component Installation	
4. Post Curing Oven	4. Electrical Inspection	
5. Trim		
6. Slot		
7. Drill		
8. Cleaning		
9. Inspection		

Each manufacturing step is modeled in SIMIO using server objects. Each object defines the processing time, the number of components that can be handled in parallel, and the associated manufacturing process. Processing times are imported from SEER-MFG, and the mapping between SEER and SIMIO is shown in Table 3. In some cases, several SIMIO steps are represented by a single SEER process; for example, activities such as cleaning and inspection are included within broader SEER manufacturing steps. Because the factory layout is notional, components are assumed to move between stations without internal transportation constraints.

Table 3. SIMIO to SEER Mapping

Manufacturing step in SIMIO	Manufacturing step in SEER
Filament Winding	Filament Winding
Heat Curing Post Curing Oven Forming Mandrel Extraction	Heat Cure
Trim	Tube Trim
Slot	Tube Slot
Drill	Tube Hole Drilling
Cleaning Inspection Electrical Inspection	Accounted for in the other manufacturing steps
Insertion of Pods into a Structure	Center Torque Box Assembly Shock Absorber Skid Installation Bulkhead Die Casting Bulkhead Machining Hoist Beam Die Casting Hoist Beam machining
Installation of Holding Alignment Pins	Missile Tube Drilling
Electrical Component Installation	Electrical Assembly
Armament Loading	Insert Missiles

Each manufacturing step also assigns a quality outcome to the part. This assignment is Boolean: A value of *True* indicates that the part passes the associated inspection, while *False* indicates failure. Parts that fail inspection are removed from the line and are assumed not to be recycled or reworked. Each manufacturing process has its own SIMIO quality-assignment routine, which samples part quality from a normal distribution, N(μ, σ). The user specifies the allowable tolerance by defining how many standard deviations from the mean are acceptable. Each step is assumed to have infinite queue capacity, so parts can accumulate while waiting to be processed; under the scenarios tested, queue lengths typically remain below 10 components. When a step has a fixed batch-processing capacity, parts wait until the required batch size is reached before processing begins. This is especially relevant for heat curing, post-curing, and cleaning. For example, if an oven can process four pods at a time, the system waits until four pods are available in the queue. In SIMIO, this batching logic is represented with a *Combiner*, which merges multiple components into a single entity for processing, and a *Separator*, which disaggregates the entity after processing. This interaction is illustrated in Figure 5.

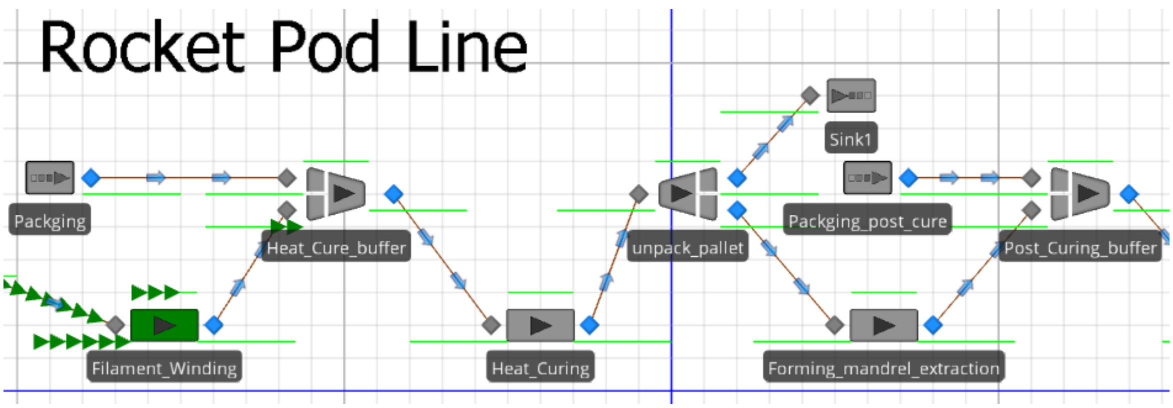

Figure 5. SIMIO Process, Combiners and Separators

Supply Deliveries

Certain manufacturing steps rely on components supplied by external manufacturers. In the current simulation, external supplies are delivered at three key stages: filament winding, electrical-component installation, and armament loading. These inputs are managed through a combination of on-hand inventory and a virtual manufacturing hub from which additional stock can be ordered.

External-component inventory is modeled using SIMIO's *Inventory* element. The element is configured with parameters that define storage capacity, resupply policy, and the resupply process. The current resupply policy is governed primarily by a threshold for replenishment and a threshold for reorder. When inventory falls to the replenishment threshold, a replenishment process is triggered to restock the warehouse.

Orders are placed with an external supplier that produces the requested quantity of goods. Although the supplier and its location relative to the manufacturing facility are not modeled explicitly, delivery time is treated as an input parameter. Users can specify nominal delivery durations as well as potential disruptions using two variables: delay frequency and a multiplier that increases the baseline delivery time when a delay occurs.

Results and Discussion

This section evaluates the simulation-driven environment developed to optimize production throughput in a notional missile-manufacturing facility. The objective is to assess the system’s responsiveness to changing operational demand through two mission-based scenarios. The first scenario assumes advance knowledge and sufficient preparation time, allowing the production floor to be adjusted to meet specified end-state inventory constraints. The second introduces uncertainty and implementation delays, highlighting the challenges of adapting late to increased demand. Both scenarios use surrogate modeling to explore design changes efficiently and identify key production bottlenecks for capacity planning and operational decision-making.

Experimental Set Up

The objective of this work is to demonstrate the modeling environment’s ability to evaluate factory-floor performance under alternative demand and resource conditions. Because the model is parametric, the number of manufacturing machines or workstations can be varied to examine the corresponding effects on throughput and production capacity. In the absence of a defined real-world layout, a notional factory floor is used as the baseline. The analysis is designed to address two questions: *(1) whether a production facility can satisfy the requirements of a given engagement scenario under a predicted expenditure profile* and *(2) how changes in available manufacturing resources influence that capability when the baseline configuration is insufficient*. To illustrate this approach, two scenarios are examined. Scenario One considers the case in which full planning time is available before mission start, while Scenario Two examines the consequences of delayed implementation and the time required for resource changes to affect production performance. Both scenarios are based on a notional factory and a notional mission. Throughout the analysis, each missile box is assumed to contain seven pods, and the number of machines on the production line is the only adjustable factor.

Surrogate Modeling

To rapidly explore how changes to the factory floor affect throughput, a surrogate model was developed based on results generated by SIMIO. This approach allows designers to easily visualize and analyze trends in performance improvements resulting from various modifications. To train the surrogate model, 2,000 cases were generated across 34 input variables, with the ranges and parameters detailed in.

Table 4

Table 4. DOE Inputs and Bounds for the 2,000 Case Run

Variable	Min	Max			
Length (in)	108	180			
Diameter (in)	8	10			
Thickness (in)	0.059	0.118			
Material	Silicon Carbide	Glass	Carbon	Armid	Aluminum Oxide
Heat Cure (Capacity)	1	10			
Cleaning Process (Capacity)	1	10			
Filament Winding (Number of Machines)	1	10			
Trimming (Number of Machines)	1	10			
Slot (Number of Machines)	1	10			
Drill (Number of Machines)	1	10			
Inspection (Number of Stations)	1	10			
Pod Insertion (Number of Machines)	1	10			
Assembly Drill (Number of Machines)	1	10			
Electrical Component Installation (Number of Stations)	1	10			
Electrical Component Inspection (Number of Stations)	1	10			
Armament Loading (Number of Machines)	1	10			
Missile Local Storage Initial (Count)	15	100			
Missile Local Storage Green (Fraction)	0.5	1			
Missile Local Storage Yellow (Fraction)	0.2	0.5			
Electrical Local Storage Initial (Count)	15	100			
Electrical Local Storage Green (Fraction)	0.5	1			
Electrical Local Storage Yellow (Fraction)	0.2	0.5			
Filament Winding Local Storage Initial (Count)	15	100			
Filament Winding Local Storage Green (Fraction)	0.5	1			
Filament Winding Local Storage Yellow (Fraction)	0.2	0.5			
Armament Delivery Delay [Normal](Hours)	5	40			
Electrical Component Delivery Delay [Normal](Hours)	5	40			
Filament Winding Delivery Delay [Normal](Hours)	5	40			
Armament Delivery Delay Extra Time of Occurrence (Hours)	0	1200			
Armament Delivery Delay [Extra](Mult)	2	10			
Electrical Component Delivery Delay Extra Time of Occurrence (Hours)	0	1200			
Electrical Component Delivery Delay [Extra](Mult)	2	10			
Filament Winding Delivery Delay Extra Time of Occurrence (Hours)	0	1200			
Filament Winding Delivery Delay [Extra](Mult)	2	10			

Scenarios and Results

Scenario One

In Scenario One, a notional mission was developed consisting of two distinct phases, each characterized by different expenditure rates defined using normal distributions. Phase One represents a high-demand period, while Phase Two reflects a lower demand. Table 5 outlines the detailed breakdown of the scenario. Demand is expressed in terms of missiles per event, and the number of missile boxes required is calculated by dividing the total missile expenditure by seven, the number of missile pods per box. For analysis, the demand data was aggregated into 25-day intervals.

Table 5. Scenario One Break-Down

Variables	Dist(mean,std)	Notes
Scenario: Continuous Battle High		
Time Interval	Normal(7,1) [hours]	Time between launch event
Number of Missiles per Event	Normal(10,3)	
Scenario: Continuous Battle Low		
Time Interval	Normal(8,1) [hours]	Time between launch event
Number of Missiles per Event	Normal(4,1)	

Assuming an initial stockpile of 1,000 missile boxes, Figure 6 illustrates the impact of demand on inventory levels. Under the given notional setup, the stockpile is fully depleted by day 425. For Scenario One, a constraint is imposed requiring that at least 400 missile boxes remain in stock by the end of the mission.

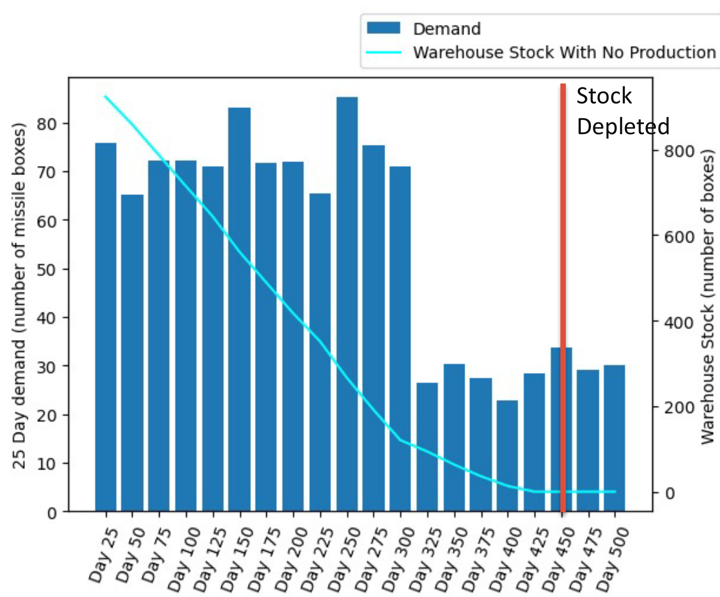

Figure 6. Demand and Warehouse Stock; No Production

As introduced earlier, a notional factory configuration was selected as the baseline. Table 6 summarizes the missile-pod parameters and the corresponding factory configuration.

Table 6. Notional Factory Set-Up

Factory Set Up		Missile Pod Params	
Heat cure (capacity)	4	Length (in)	152.37
Post Heat Cure (capacity)	4	Diameter (in)	9.43
Cleaning process (capacity)	2	Thickness (in)	0.1
Filament Winding (machine #)	1	Material	Carbon
Mandrel Extraction (machine #)	1		
Trimming (machine #)	2		
Slot (machine #)	5		
Drill (machine #)	6		
Inspecting (station #)	4		
Pod insertion (machine #)	1		
Drill assembly (machine #)	5		
Electrical components install (machine #)	7		
Assembly inspection (station #)	1		
Armament loading (machine #)	4		

Figure 7 shows the performance of the notional factory configuration under the mission described above. This configuration yields a throughput of 6.2 missile boxes per 25-day period.

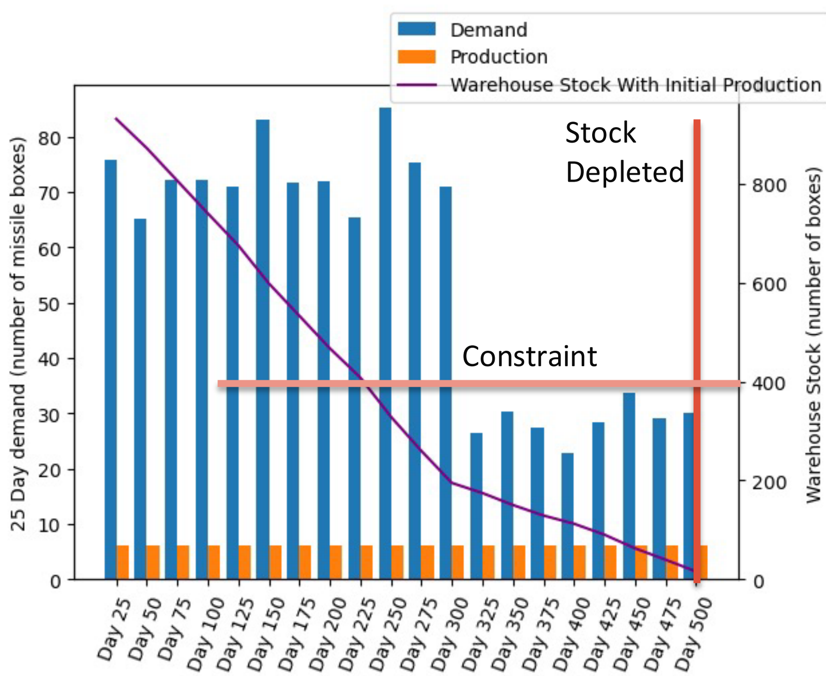

Figure 7. Demand, Production, and Stock Under a Notional Factory Setup

While this setup extends the stockpile to day 500, it fails to meet the requirement of maintaining at least 400 missile boxes in inventory by the end of the mission. To satisfy this constraint, adjustments to the production floor are necessary. The first step is to determine the required throughput, which can be calculated by comparing total inflows and outflows, as shown in Equation 1. In this equation, *P* represents the desired throughput, *sg* is the number of 25-day segments, S_i is the starting stock, S_L is the target ending stock, and D_k denotes the demand. For this scenario, the required throughput is 25.4 missile boxes per 25-day period, equivalent to 101.6 missile boxes per 100 days.

$$P = \frac{S_i - S_L - \sum_{k=1}^{sg} D_k}{-20} \tag{1}$$

The objective is to increase production from 6.2 to 25.4 missile boxes over a 25-day period (equivalent to 101.6 boxes over 100 days). To identify the production steps that most strongly affect throughput, the simulation results were analyzed using JMP's predictor-screening capability. As shown in Figure 8, the most influential steps are filament winding, pod drilling, pod trimming, and heat curing.

Predictor	Missile box made Response Contribution	Missile box made Response Portion	Missile box made Response Rank	Missile pods made Response Contribution	Missile pods made Response Portion	Missile pods made Response Rank
Filament Winding Number of Machines	915999	0.2612	1	48767684	0.2553	1
Drill machines Number of Machines	761031	0.2170	2	41384490	0.2166	2
Trimming Machines Number of Machines	748562	0.2135	3	40306383	0.2110	3
Mandrel Extraction Number of Machines	681230	0.1943	4	39552679	0.2070	4
Heat cure capacity	194383	0.0554	5	10002908	0.0524	5
Post Heat cure capacity	164301	0.0469	6	9438305.9	0.0494	6
Filament Winding time (min)	5902	0.0017	7	422825.81	0.0022	7
Insertion Pods Structure time (min)	5737	0.0016	8	350097.68	0.0018	8
Control: Armament Delivery Disruption Amount of Delay (Time)	4460	0.0013	9	6930.6843	0.0000	34

Figure 8. JMP Predictor Screening, Missile Box and Missile Pods Throughput

Using the surrogate model, a profiler was generated to show how missile-box throughput changes in response to modifications on the factory floor. The profiler reports throughput in 100-day intervals; values for a 25-day period can therefore be obtained by dividing by four. As shown in Figure 9, increasing the number of filament winding machines produces a corresponding increase in overall system throughput.

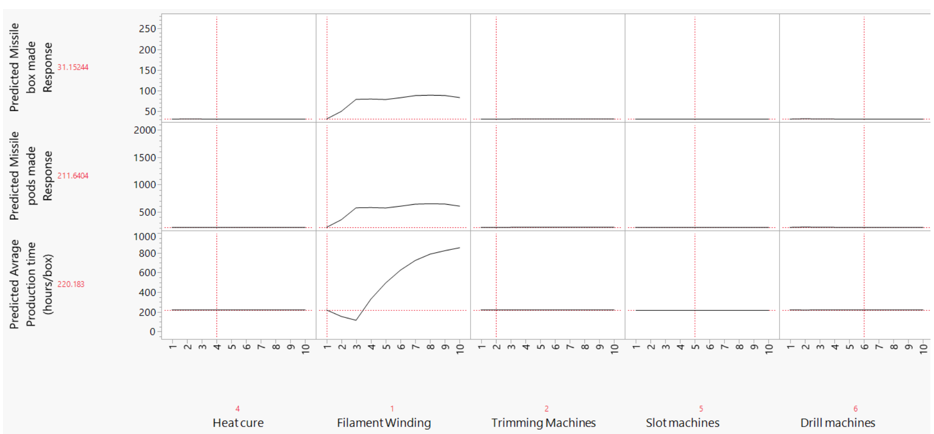

Figure 9. Predictor Profiler With Initial Setup

Increasing the number of filament-winding machines from one to three raises throughput from 31 to 78.8 missile boxes per 100 days. Although this is a substantial improvement, it remains below the required 101.6 boxes per 100 days. No single change is sufficient to close the gap. However, the predictor-screening results in Figure 8 indicate that drilling and trimming are also important contributors to throughput. The updated facility layout therefore adds three filament-winding machines and one trimming machine. As shown in Figure 10, this configuration increases throughput to 103.6 missile boxes per 100 days, or 25.9 boxes per 25 days. That output exceeds the target throughput of 25.4 missile boxes per 25 days. Figure 11 shows the resulting relationship between stock levels and demand under the revised setup.

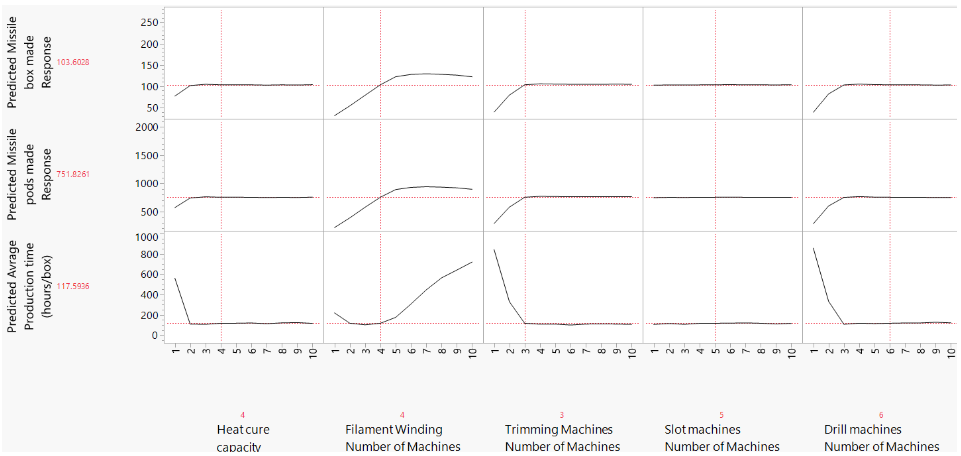

Figure 10. Predictor Profiler With the Number of Filament-Winding Machines Increased by Three and Trimming Machines by One

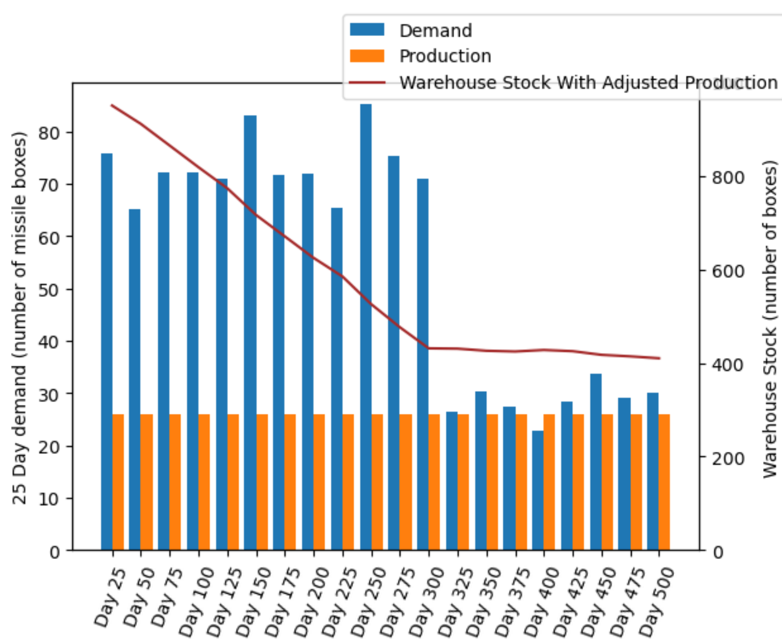

Figure 11. Demand, Production, and Stock Under the Adjusted Factory Setup

Scenario Two

Scenario Two begins where Scenario One ends. After the initial mission concludes, it is determined that an additional surge in missile expenditure is required. This introduces a new 300-day period of high demand. The scenario begins with a remaining stockpile of 410 missile boxes, which is the ending inventory from Scenario One. Unlike the previous case, this surge is not identified until day 500, so production changes cannot be implemented in advance.

In Scenario Two, modifications to the production floor require time to implement. For example, if a set of changes takes 100 days, the higher throughput is not realized until 100 days into the scenario. Each machine change is assumed to require 25 days, and the time required for multiple changes is additive; for instance, two changes require 50 days. In addition, all changes must be completed before the new production capacity becomes active, so staggered implementation is not allowed.

As shown in Figure 12, the factory configuration established in Scenario One is insufficient to meet the new demand, and the stockpile is fully depleted by day 725.

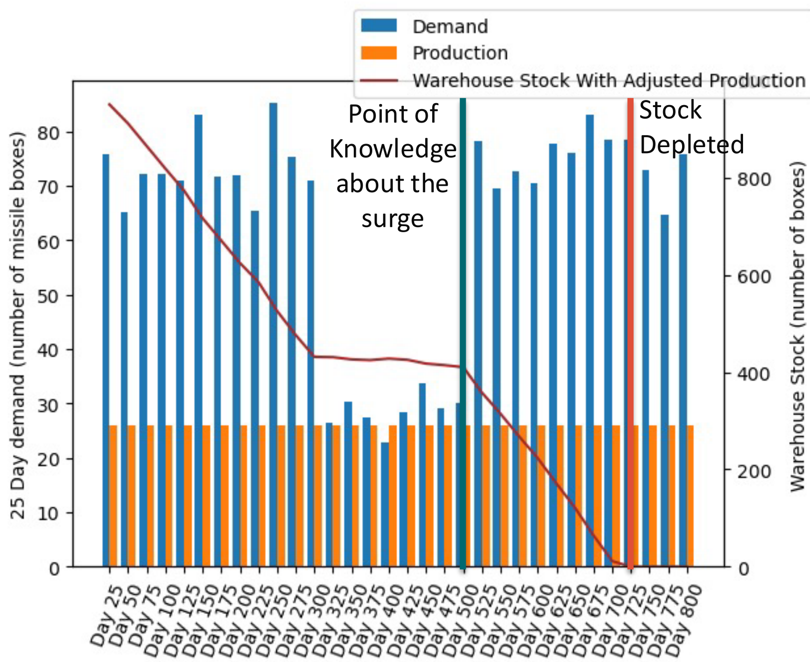

Figure 12. The Extended Scenario With Throughput Given by the Adjusted Manufacturing Facility

Given these constraints, an equation can be formulated that defines the relationship between the number of changes and the required throughput. This relationship is expressed in Equation 2, where *sg* represents the total number of 25-day segments, *c* is the number of 25-day segments at new production rate, P_1 is the initial throughput, and P_2 is the modified throughput.

$$c(P_1 - P_2) < S_i + (sg)P_1 - \sum_{k=1}^{sg} D_k \qquad (2)$$

The result of this relationship is shown in Figure 13. Since each change takes 25 days to implement, adding a single change reduces the time available at the new throughput rate from 12 segments to 11. Initially, the increase in required throughput is approximately linear, between

8 and 12 segments under the new production rate. Beyond this range, however, the required throughput begins to rise much more sharply.

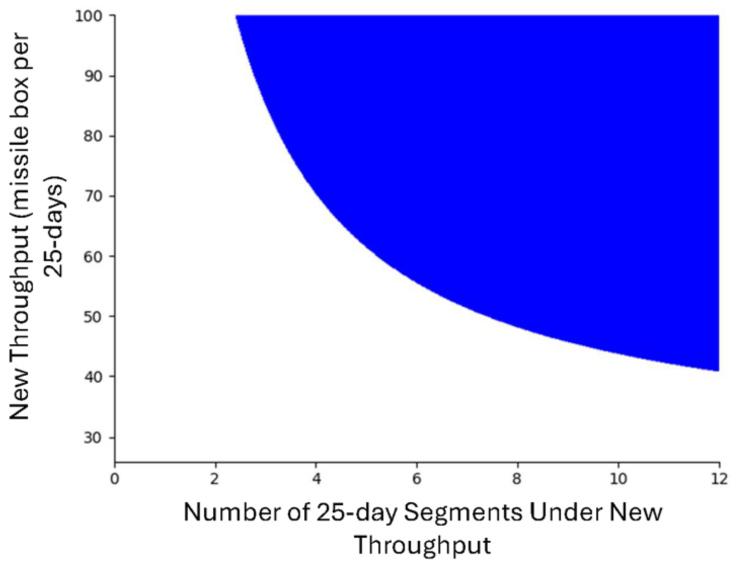

Figure 13. Trade-Off Between the Number of Changes and the Required New Throughput

As the number of changes increases, so does the required throughput. By examining the throughput deficit using Equation 3, it is possible to assess whether the new mission demand can be met.

$$Deficit = \frac{New\ Throughput - Needed\ Througput}{Needed\ Througput} \tag{3}$$

The results, shown in Table 7, indicate that up to four changes reduce the throughput deficit. When the number of changes increases to six, requiring 150 days to implement, the deficit begins to worsen. This occurs because the throughput gains from adding machines are outpaced by the increase in required production. As illustrated in Figure 14, implementing four changes allows the stockpile to last until day 750. Given the constraints of this scenario, reaching day 750 is the best achievable outcome.

Table 7. The Change of Throughput Deficit Based on the Number of Changes

Number of Changes	Throughput Deficit (%)
1	27 %
3	15.6 %
4	13.76 %
6	15.27%

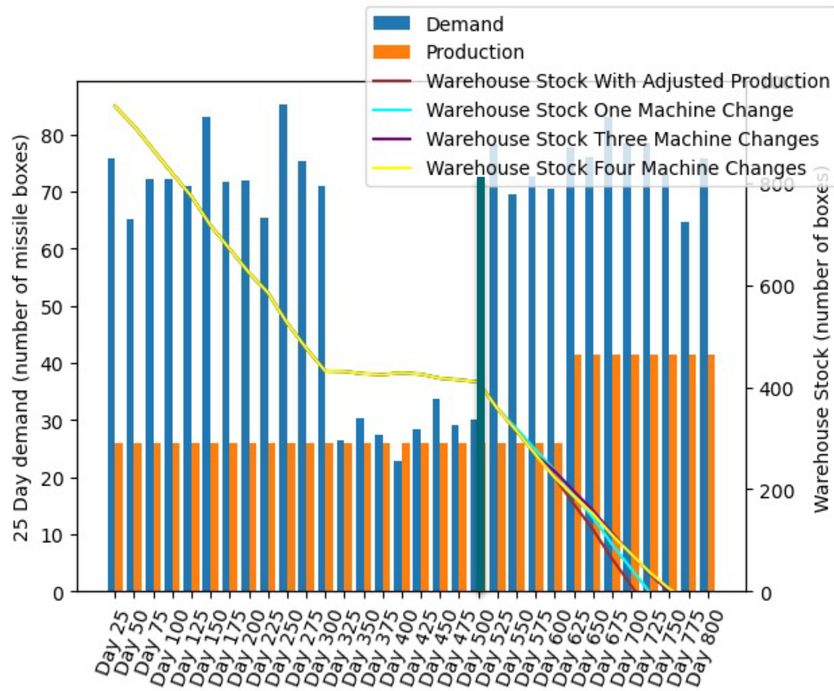

Figure 14. The Extended Scenario With Throughputs With Different Number of Changes

Conclusions and Future Work

A comprehensive parametric modeling and simulation environment was developed to examine how manufacturing capacity affects the ability of the broader factory-to-field system to satisfy operational demand. Using SEER-MFG to estimate manufacturing-process times and SIMIO to quantify throughput for a given factory configuration, the environment provides a rapid and flexible means of evaluating how changes on the factory floor influence production performance. In total, the model includes 65 adjustable inputs, allowing users to examine the effects of assumptions related to quality control, machine quantity and capacity, delivery delays, and inventory-storage limits. In this sense, the contribution of this work is a manufacturing-focused component of a larger decision-support framework intended to connect demand, logistics, and production.

Two operational scenarios were used to demonstrate these capabilities. In the first scenario, all production changes were assumed to be implementable before mission start. The results showed that the environment can identify whether a given factory configuration is sufficient to satisfy mission demand and can rapidly evaluate how changes in production resources affect throughput. In the second scenario, production changes had to be made during mission execution, introducing time-dependent implementation constraints. This case demonstrated the environment's value for assessing the consequences of delayed action and for characterizing the relationship between implementation time and the level of production capacity required to recover. Taken together, these scenarios show that the modeling environment can support timely and structured assessment of manufacturing responsiveness under both preplanned and dynamically evolving conditions.

More broadly, this work addresses a critical need in defense planning: the need to link operational demand to the logistics, manufacturing, and production systems required to sustain combat capability. In contested or time-sensitive environments, it is not enough to understand demand in isolation; planners must also understand whether the industrial base and supporting supply chain can respond at the speed and scale required by the mission. Studies that connect demand forecasting, logistics timelines, manufacturing constraints, and production throughput are therefore essential for evaluating resilience, identifying bottlenecks, and informing resource-allocation decisions across the end-to-end system. The manufacturing and production models presented here contribute to that larger objective by providing a structured way to assess how factory-floor decisions influence the ability to meet scenario-driven demand.

Future work should focus on increasing both the fidelity of the manufacturing representation and the degree of integration across the broader modeling environment. Important extensions include relaxing simplifying assumptions and incorporating additional decision variables such as workforce availability, work-shift structures, material characteristics, part geometries, maintenance constraints, and facility-layout considerations. Future studies should also examine how existing manufacturing facilities and supply chains absorb new or surge demand without degrading support to other high-priority operations. Most importantly, the manufacturing model should be fully integrated with the Phase I logistics model described in the Methodology and associated with Figure 1 to enable a much-needed end-to-end analysis of how operational demand propagates through transportation, inventory, and production systems. Such integration would significantly improve the utility of the framework as a decision-support tool for planning resilient defense industrial operations under uncertainty.

Acknowledgments

This research was funded by Northrop Grumman Corporation.

SEER-MFG is a trademark of Galorath Incorporated.

SIMIO is a trademark of Simio LLC.

MLRS is a trademark of Lockheed Martin Corporation.

References

AIAA Digital Engineering Integration Committee et al. (2020). *Digital twin: Definition & value* [AIAA and AIA Position paper].

Battaïa, O. (2013). A taxonomy of line balancing problems and their solution approaches. *International Journal of Production Economics*, 259–277.

Birbasov, N. S., Wang, P., Cimtalay, S., Pinon-Fischer, O. J., Balchanos, M. G., & Mavris, D. (2025). *A modeling and simulation environment for the quantitative assessment of operational and technological improvements in military logistics*. AIAA SCITECH 2025 Forum, Orlando.

Cantrell, S., Margolis, C., Krauss, M. F., & Mavris, D. (2025). Digital twins for sustainment-oriented wargaming. *AIAA SCITECH 2025 Forum.* AIAA.

Falkenauer, E. (2005). Line balancing in the real world. *Proceedings of the International Conference on Product Lifecycle Management PLM*, 360–370.

Galorath. (n.d.). *SEER for manufacturing (SEER-MFG).* Retrieved May 2025 from https://galorath.com/cost-estimation/seer-mfg-manufacturing/

Hoem, M. T., Jurasek, G. M., Manuel, L. J., Samal, S., Kampezidou, S. I., Jonchay, T. S. . . . Dunn, K. (2025). A quantitative assessment of the environmental and financial impact of digital twins on large scale aerospace logistics operations. *AIAA SCITECH 2025 Forum.* AIAA.

Libby, S., Siedlak, D. J., Solano, H. D., Pinon Fischer, O. J., & Mavris, D. N. (2017). Cost-capability analysis of UAS family and flexible factory design. *17th AIAA Aviation Technology, Integration, and Operations Conference.* AIAA.

Ling-Temco-Vought Aerospace and Defense. (n.d.). *Making of the Multiple Launch Rocket System (MLRS).* Retrieved from https://www.youtube.com/watch?v=NYRNjghl61U

Metcalf, J. G., & Laffey, J. A. (2023). *Integrating digital twin concepts to enhance agility of the United States Marine Corps' decision support framework.* Naval Postgraduate School.

Misra, M., Leicht, B., Vallejos, J. P., Green, N., Ozcan, M., Pinon Fischer, O., & D.N, M. (2025). Integration of engine digital twin in contested logistics framework. *SCITECH.*

National Defense Industrial Association. (2025). *NDIA & DRIVE contested logistics workshops summary & recommendations.* NDIA.

Paksoy, T. Ö. (2012). Supply chain optimisation with assembly line balancing. *International Journal of Production Research*, 3115–3136.

Pinon Fischer, O. J., Chu, S., Huyn, D., dos Santos, M., Cox, A., Winecoff, D., & Waliliong, A. (2022). *Model-based systems engineering for digital manufacturing: A proof-of-concept.* MxD. https://apps.dtic.mil/sti/pdfs/AD1182010.pdf

Rahman, S. A. (2023). A simulation-based approach for line balancing under demand uncertainty in production environment. *Winter Simulation Conference (WSC).*

Schlichting, G. S., Reddy, A., Murphy, D., Oroz, J., Pinon Fischer, O. J., & Mavris, D. N. (2024). Contested logistics operating under digital support. *AIAA SCITECH.*https://www.doi.org/10.2514/6.2024-0481

Shin, S. Y. (2017). Discrete-event simulation and integer linear programming for constraint-aware resource scheduling. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 1578–1593.

Siedlak, D. J., Schlais, P. R., Pinon, O. J., & Mavris, D. N. (2015). Supporting affordability-based design decisions in the presence of demand variability. *International Manufacturing Science and Engineering Conference.* ASME.

Sieldak, D. J., Pinon, O. J., Schlais, P. R., Schmidt, T. M., & Mavris, D. N. (2018). A digital thread approach to support manufacturing-influenced conceptual aircraft design. *Research in Engineering Design, 29*(2), 285–308.

Wang, T., & Liu, Z. (2021). Digital Twin and its application for the mainenance of aircraft. In *Handbook of nondestructive evaluation 4.0* (pp. 1–19).

PANEL 17. BEYOND THE VALLEY OF DEATH: VENTURE STUDIOS AND STRATEGIC SMALL BUSINESS GROWTH

Thursday, May 7, 2026, 1145 – 1300 ET (0845 - 1000 PT)

Panel Summary:

Bridging the gap between a laboratory concept and a fielded military capability requires a fundamental shift in how the Department of War (DoW) incentivizes small businesses and simulates policy impacts. This panel introduces transformative models for scaling innovation, including the Small Business Engineering Resource (SBER), which pivots the SBIR program from perpetual research to assured engineering execution. Complementing this are cutting-edge tools like the Agentic AI-based "Policy Test Lab," designed to simulate acquisition reforms at machine speed, and the Venture Studio model, which applies proven commercial scaling methodologies to accelerate the development of operational warfighting capabilities.

Chair: Pete Modigliani, CEO, Creative Defense Network (CDN)

Panel Presenters:

Small Business Engineering Resource: Breaking the SBIR-Mill Business Model and Creating Big Companies from Small Ones – *The Honorable Nickolas Guertin, Virginia Tech National Security Institute*

Venture Studios to Scale Warfighting Concepts to Capabilities – *Dr. Peter Khooshabeh, DEVCOM Army Research Laboratory (ARL)*

The Policy Test Lab: An Agentic AI-based Simulation Tool – *Dr. Douglas J. Buettner, Acquisition Innovation Research Center (AIRC)*

Pete Modigliani— has held key positions in prestigious organizations like Govini, Atlantic Council, Beacon Global Strategies LLC, MITRE, Alion Science and Technology, and more. He also runs the Defense Tech and Acquisition Substack, where he shares valuable insights and updates on defense acquisition.

Pete led strategic initiatives and reforms across the Department of Defense and Intelligence Community acquisition enterprises to enable delivery of better capabilities faster. He led MITRE's efforts within the Office of Secretary of Defense (A&S) for software acquisition, portfolio management, and acquisition enablers. He advised acquisition executives and program managers on transforming their strategies, initiatives, and processes.

He co-founded the Acquisition in the Digital Age (AiDA) platform to disrupt and digitally transform the acquisition enterprise. It led to the development of the transformational Adaptive Acquisition Framework (AAF) for USD(A&S) and the acquisition workforce. This included developing and shaping the new Congressionally directed acquisition pathways for Middle Tier of Acquisition and Software Acquisition.

He previously was a part of the Congressionally directed Section 809 Panel where he championed strategic portfolio management, simplifying acquisition, and requirements reforms.

Pete served as an Air Force officer and program manager within the Joint STARS and Mission Planning programs. He led the strategy development of a billion-dollar enterprise contract, modernization programs, and managed multiple smaller acquisition and systems engineering efforts. While serving in this capacity, he partnered with industry executives to prototype a B2G portal to foster collaboration on Air Force capability needs and industry offerings.

He began his career as an industrial engineer with Xerox, 3M, and ITT Automotive. He received a bachelor's degree in industrial engineering from Rochester Institute of Technology and a Master of Business Administration from Boston College.

Small Business Engineering Resource: Breaking the SBIR-Mill Business Model and Creating Big Companies from Small Ones

Nickolas H. Guertin—has 30+ years of leadership in defense acquisition and testing. He served as Assistant Secretary of the Navy for Research, Development, and Acquisition, overseeing a $130 billion portfolio and 130,000 personnel. Before that he served as the Pentagon's Director of Operational Test and Evaluation. A retired Navy civilian and submariner, he is recognized nationally for advancing acquisition strategy, open-systems architecture, and operational testing, and continues shaping defense innovation through research and advisory roles.

Howard Reichel—serves as Senior Vice President and Chief Operating Officer at In-Depth Engineering. In-Depth is a small business provider of combat systems solutions, waterfront support, adaptive e-learning and AI/ML solutions. In-Depth has a rich combat system development legacy - designing and delivering real-time mission critical "weapons-safe" software systems to the Department of the Navy. Solutions include geo-spatial visualization and fusion products, tactical-control and weapons-control solutions for heavyweight, lightweight and anti-torpedo torpedoes, real-time image rendering and augmentation solutions, tactical decision aids, algorithms and estimation theory, and system infrastructure products. Our ongoing applied and advanced research and technology programs drive the state of the art and lay the groundwork for next generation warfighting capability. Prior to his position at IDE, Mr. Reichel served as a member of the Senior Executive Service at the Department of Homeland Security and as a Program Manager within PEO IWS.

Abstract

The Small Business Innovation Research (SBIR) program funds important applied research that often fails to transition into fielded military capability. Recent congressional activity highlights continued frustration with this gap (Ernst, 2026). This paper argues that pending SBIR reauthorization (as of this writing) does not address the root cause: Topics, statutory funding caps, and governance remain misaligned with acquisition pathways and current acquisition transformation efforts (Bresler, 2023; Department of Defense, 2025). We redefine the pervasive "SBIR-mill" outcome not as a moral failing of firms, but as a predictable result of government-created incentives that prioritize recurring research over production at scale, inhibit acquisition success, and prevent graduation from the program (Bryant, 2025).

To rectify this, we propose recasting SBIR as the Small Business Engineering Resource (SBER): a program oriented toward engineering execution, qualification, and integration so solutions are mature enough for adoption and scaling. This is not a cosmetic rebrand; it is a shift from early-stage R&D toward deliverables that are programmatically and financially embedded in Programs of Record, so transition is planned and funded from the start.

We outline a research design to evaluate two measurable transition outcomes: (a) integration into a Program of Record under a prime contractor or (b) displacement of an incumbent supplier enabled by open interfaces and competed technology insertion. We derive testable hypotheses and a data plan using federal acquisition data, with policy levers including fewer, larger awards (e.g., >$20 million), topics tied to Program of Record requirements with accountable PM sponsorship, and budgeting aligned to PPBE with explicit funds for integration, test, qualification, delivery, and iterative improvement.

The Structural Inadequacy of the SBIR Program for Scaling Capability

The Persistent Challenge of the "Valley of Death" in Defense Acquisition

The SBIR program was established with a dual mandate: stimulate innovation among small businesses and ensure resulting technologies transition into federal acquisition programs (S. 881, 1982). While the program succeeds at funding early-stage R&D, it repeatedly fails at transition and firm growth in defense—often described as the "Valley of Death" (Bresler, 2023). That transition failure is a structural deficiency in the current acquisition model and supporting legislation.

SBIR Phase I–II efforts typically mature technologies to roughly TRL 4–6, but fielding requires "productionization," qualification, test, systems engineering, and integration into existing platforms. Those activities are necessary for TRL 7–9 outcomes yet are largely unsupported by SBIR's typical funding levels and timelines, leaving small firms without a credible path to sell into Programs of Record (PORs) or through primes (Borda, 2025; McNamara, 2025) and are unattractive for possible commercialization outside of the military market.

The Phase II Funding Cliff and Financial Incoherence

The core flaw in the current SBIR structure is the financial incoherence surrounding the transition phases. Statutory caps restrict Phase II awards, which are intended for continued research and prototyping, to approximately $2.1 million over 15–24 months, though waivers can sometimes increase this amount. This ceiling is generally adequate for advanced prototype development (TRL 4–6) but is fundamentally insufficient to fund the necessary qualification, testing, and systems engineering (SE) required to achieve TRL 7–9 readiness for major acquisition programs. The institutional and budgetary context confirms that typical SBIR scope, such as $2 million over three years, rarely transition because they have insufficient maturity to bridge the gap to higher acquisition phases.

These difficulties are exacerbated by PPBE timing and color-of-money constraints. Phase II is usually funded with 6.2/6.3 RDT&E, while the work needed to field capability (advanced development, integration, qualification, and initial production planning) must be funded on different timelines and often from different appropriations. The SBIR assumption that a successful Phase II "ROI event" will trigger timely follow-on funding is inconsistent with PPBE, creating a predictable multi-year gap and reinforcing the Valley of Death (Gallo, 2025).

Institutional Barriers and Program Manager Incentives

The cultural landscape of defense acquisition also contributes significantly to transition failure. Program Managers (PMs) operating large, complex PORs must balance execution risk and avoid immaturity in the delivered product, which may also result in operational risk to the fleet and force. As such, their careers and the success of their programs depend on mitigating schedule slips and technical failures for the validated and funded requirements. Integrating innovative technology from a small, non-traditional vendor is often viewed as creating more difficulty and programmatic burden that introduces challenges to implementation into an already difficult process. Limited Phase II funding, resulting in early prototypes from Phase II efforts exacerbate the concern.

PMs also face weak incentives to absorb integration risk from non-traditional vendors, and major program budget and milestone timelines rarely align to SBIR periods of performance. Any effective SBIR restructuring therefore must treat transition efforts as "mini acquisition programs" with integrated planning, budgeting, and governance—not stand-alone research projects.

A related structural issue is topic creation: Many SBIR topics originate in S&T communities and are not tied to validated, funded POR requirements. When topics are disconnected from scheduled requirements, even a strong technical outcome can be doomed by design because no funded transition path exists.

Defining the "SBIR Mill" Phenomenon: Blaming the Victim and Structural Incentive Failure

Critique of the SBIR program often centers on the existence of firms that perpetually win Phase I and Phase II awards while rarely (or never) transitioning the resulting technology to commercial or defense utility; these are the so-called "SBIR Mills." The conventional analysis frequently characterizes this as a moral failing or inefficiency on the part of the businesses themselves (VanRoo, 2022). We make no attempt here to attack the firms. We assert instead that they are responding to the business environment and incentive structures they must operate in. Deeper analysis reveals the lack of break-out companies that take the SBIR path and ultimately succeed from being coupled to, and subjugated by, the defense industrial ecosystem. This is fundamentally a structural incentive failure set up by the government's legislation and associated policy implementation, not an ethical or competence deficiency among the participating firms.

At the conclusion of Phase II, many firms face no viable funded path to transition. In that environment, firms rationally pursue additional Phase I/II awards to retain specialized staff and remain solvent, reinforcing a cycle of repeated research awards ("the mill") rather than scale-up and production. Structural changes, especially acquisition alignment, funded transition pathways, and enforceable governance—are therefore *required* to replace the mill dynamic with predictable transition.

Policy Effectiveness and Systemic Constraints

Recent legislative proposals to reauthorize SBIR seek to address longstanding concerns regarding commercialization and transition. The new language contained in the pending SBIR Reauthorization Act (S. 3971) introduces several meaningful changes, including: (1) a "strategic breakthrough" Phase II authority of up to $30 million over four years, (2) a requirement for matching funds, (3) a commitment for inclusion in the Program Objective Memorandum (POM) at the Program Executive Office level or higher, and (4) expanded authorities for Phase III and improved data tracking.

These provisions acknowledge the transition problem, but they create a selective pathway rather than restructuring the baseline program.

Past legislative efforts, such as the SBIR and STTR Extension Act of 2022, added transition standards for Multiple Award Winners (MAWs) intended to filter out "mills." Results to date suggest limited practical effect: Only a small number of businesses approach the imposed limits, and few face meaningful consequences (Naval Submarine League, 2012).

These efforts suggest that tighter participation rules alone do not solve the transition problem. Proposals such as the INNOVATE Act (S. 853) emphasize award caps and submission limits but do not fix the structural requirements for transition: aligned topics, appropriate funding, executable contracts, and governance that ties prototypes to acquisition pathways and competition mechanisms (GAO, 2025).

Policy Proposal: SBER and the Dual-Path Transition Mandate

To rectify this deep structural misalignment, we propose the SBIR program be recast as the Small Business Engineering Resource (SBER) Program. This proposal creates a strategic shift in focus from the execution of early-stage R&D project at TRL 4–6 to the execution of advanced development engineering. With a goal of getting capability into the hands of operators/warfighters, we reject the siren-song of driving immature products into the field. This popular, but ineffective notion, ultimately shifts acquisition simplicity into operational risk. These innovations must be built to minimum viable capabilities that include:

- flexibly/modular designs that fluidly accept user feedback into successive development,
- excellence in integration into the target environment,
- sufficient automated testing that ensures the thing will work when called upon at times when lives matter,
- sufficient qualification for use on the battlefield, and
- excellent field training to ensure the user knows what right looks like.

This would take the prototype up through TRL 7–9 to be considered ready for initial deployment to the field. SBER would be designed, funded, and governed in a manner like the Navy's Undersea Warfare 6.3 Advanced Processing/Capability Build (APB/ACB) programs, which has a decades-long track record of using small businesses to deliver innovative yet mature R&D solutions directly into PORs (Guertin & Kaliz, 2005). Small business solutions are executed with sufficient funding and governance to resemble miniature acquisition projects. There is no valley of death in those R&D programs, as the scope of work is important, the viability of the technology solution is known, and the small business, large business and oversight activities plan the process and are funded for success from the beginning.

Structural Mandates for SBER

The shift to SBER necessitates specific structural mandates to ensure that technologies are anchored to real lethality needs from inception:

1. Mandatory Topic Sponsorship: SBER Topics must be approved at a sufficiently senior level—specifically by a responsible/accountable Program Manager (PM) and a Service Chief Warfare Sponsor Resource Officer (RO). This governance change takes the topic creation and funding guesswork out of the hands of people who are not responsible for outcomes and guarantees that the SBER efforts are aligned with funded, scheduled requirements within a specific POR.
2. Fewer, Larger Awards: The traditional Phase II limits must be redesigned by allocating fewer, but significantly larger, awards, targeting $10 million to $20 million per effort. This concept is partially echoed in current legislative proposals, such as the INNOVATE Act, which proposes 'strategic breakthrough' grants or Phase III defense-specific awards up to $30 million to push commercialization and fielding. The new SBER award structure mandates that these funds span the entire prototype lifecycle, through qualification, integration, test, and initial delivery of a minimum viable capability, addressing the previously identified deficits in 6.4/6.5 RDT&E funding.
3. Execution through the POR: The SBER award execution must be embedded within and managed by the POR. This would typically involve the small business acting as a subcontractor to the prime contractor or as a displacing capability as described below.

The Dual-Path Transition Model

Central to the SBER concept is the establishment of two explicit, measurable, and structurally enforced transition pathways that define success not merely as technology development, but as system insertion: Path A (Subcontract Integration) and Path B (Incumbent Displacement). These pathways provide the foundational policy framework for success measurement and strategic investment.

This structural shift requires a clear understanding of the investment mechanics and objectives, as outlined in the following comparative table.

Table 1. Modeling the Structural Shift: Current SBIR vs. Proposed SBER Award

	Current SBIR Phase II (Statutory Cap)	**Proposed SBER Transition Award ($10–$20M)**	**Policy Objective Alignment**
Primary Focus	Continued Research and Prototype (TRL 4–6)	Engineering, Qualification, and Integration (TRL 7–9)	Mitigate TRL/Acquisition Gap
Funding Cap (Typical/Statutory)	$2.1 Million	$10 Million – $20 Million (Single Award)	Fund Test/Systems Engineering and Production Scale-up
Duration (Typical)	15 to 24 Months	3–5 Years (Aligned to POR Milestones)	Match POR Schedules
Color of Money Target	RDT&E 6.2/6.3	Integrated RDT&E 6.4/6.5 and Procurement Planning	Fund Prime, Small Business and Oversight Activities. Bridge Phase Gaps
Transition Authority	Uncertain Phase III, often limited by traditional FAR-based contracts	Assured Phase III via subcontract to Prime or MTA/CSO/OTA Bridge Contracts	Guarantee Production Path

Conceptual Model: The Dual-Path Transition Framework (SBER)

The SBER model institutionalizes two transition strategies intended to ensure adoption within the acquisition lifecycle while growing small businesses. Both paths leverage the DoD's Modular Open Systems Approach (MOSA) as a key enabler for integration and competition.

Path A: Subcontract Integration (SB-as-Module)

Path A focuses on the seamless integration of the SBER-developed module into an existing or planned POR structure. This path transforms the small business performer into a structured and supported subcontractor to the prime contractor responsible for the overall system integration.

The efficacy of Path A is driven by mandatory governance changes. This will require both the government PM and RO joint approval to ensure that the technology is fully integrated into the programmatic, funding, and technical structure and a critical component of the POR. This top-down mandate overcomes the PM's inherent resistance to risk, as the technology is organizationally and financially supported from the highest levels of acquisition authority and operational commands. For this path, the prime contractor serves as the key integrator and sub-contracting manager. By embedding the SBER execution directly through the POR, the structural barriers that typically isolate small businesses from major weapon systems are systematically dismantled.

Path A also uses the SBER award as a contractual incentive for the prime contractor to integrate the small business module. By funding the integration work and making the module a required program element, the government changes the prime's cost/benefit calculus and shifts initial integration risk away from the small business.

Path A relies on open interfaces (or the mandated creation of MOSA interfaces) to reduce integration burden. Success can be measured with operationally meaningful metrics such as fielded units using the module, integration cycle time, and post-fielding defect/deficiency rates.

Path B: Incumbent Displacement (SB-as-Challenger)

Path B is the structural enforcement mechanism designed to promote competition and technological refresh by allowing the SBER small business to function as a challenger, displacing incumbent suppliers who have high cost, perform poorly, or rely on proprietary/closed architectures.

Path B depends on enforceable open standards: the government must mandate MOSA compliance in major programs and fund the reference architectures and conformance testing needed to make “plug-and-play” competition credible. This reduces vendor lock and lowers selection risk for PMs considering a challenger solution.

To make displacement a reality, Path B demands significant government investment in reference architectures and conformance test environments. These resources are serving as objective testing laboratories where the small business can rigorously prove its fluid integration (a.k.a. “plug-and-play”) conformance to the open standard. By funding and stewarding these open interfaces and the necessary testing infrastructure, the government manages systemic risk and avoids future vendor lock that typically plagues proprietary systems. This approach effectively standardizes the component market within a given architecture, shifting competition away from specific vendors that control interfaces into those that provide the best technical solution (cost, performance, agility). This reduction in technological risk is necessary to allow PMs to confidently select a small business solution over a familiar incumbent.

Success in Path B can be measured by displaced modules, cost/schedule/performance differences versus incumbents, and insertion cycle time. Enabling these outcomes requires clear data-rights terms and planned technology-insertion windows so competitions can be executed on predictable schedules.

Table 2. Operational Mechanics of the Dual-Path SBER Transition Model

Transition Path	**Path A: Subcontract Integration (SB-as-Module)**	**Path B: Incumbent Displacement (SB-as-Challenger)**	**Critical Policy Levers**
Acquisition Goal	Seamlessly integrate capability into a POR structure.	Institutionalize market competition and technological refresh.	Governance & Standards
Technical Enabler	Existing or newly established Modular Open Systems Approach (MOSA) interface.	Government-stewarded Open Interfaces and Funded Conformance Benches.	MOSA Mandate Enforcement
Required Governance	Mandatory O-6/Pentagon Sponsorship; Execution via POR.	Standardized IP/Data Rights; Publicly published Tech-Insertion Windows.	Contractual Clarity
Metrics of Success	Fielded units, integration cycle time, defect escape rate.	Displaced modules, cost/performance delta vs. incumbent, insertion cycle time.	Quantifiable Outcomes

Putting the Two-Path Concept Into Motion – Governance and Scaling

The transformation of SBIR into SBER is dependent on radically restructuring the funding and governance architecture to treat the small business as an essential resource for production rather than a source of exploratory research.

Financial Mandates: Justification for Fewer, Larger Awards – and Alignment With PPBE

Fewer, larger SBER awards (roughly $10–$20 million) would cover the cost of qualification and effort needed to address sufficient supportability to move from a TRL 6 prototype to an operationally suitable TRL 9 capability. The current Phase II cap is generally insufficient to fund the integration, test, certification, and production-engineering work needed for scale and operational impact (DoD, 2024).

Funding this full transition scope changes the PM risk calculus from "will the prototype work?" to "does the qualified component meet the interface and performance requirements?" Compared to reforms that primarily expand entry pathways (e.g., Phase IA concepts), the SBER approach targets the system bottleneck: resourcing acquisition readiness and integration.

When an SBER module is planned and executed within the PPBE system through certification, integration, test, and delivery, the program can reduce both acquisition and operational risk while enabling faster, more predictable insertion.

In today's defense innovation climate, private capital is increasingly directed toward *dual-use* companies that can scale in commercial markets while also delivering military capability (e.g., venture, growth equity, and corporate strategic investment aligned to national security). Dual-use is attractive because commercial demand can finance manufacturing readiness, talent, and iterative product improvement, while DoD demand can validate requirements and provide resilient revenue. But private capital is not a substitute for an executable government transition path: Without clear demand signals (e.g., POR sponsorship), sound intellectual property strategies, and planned insertion windows, investors will discount DoD revenue as uncertain and are more likely to push firms to prioritize purely commercial roadmaps or repeated R&D awards. The SBER model improves dual-use opportunities and attractiveness for investment by turning transition into a governed, budget-aligned engineering and integration effort with measurable outcomes, allowing external capital to complement (rather than compensate for) PPBE and acquisition misalignment.

Addressing Logistics and Classification Barriers

Beyond direct financial investment, SBER must also overcome logistical barriers to integration. Small businesses frequently lack the facility clearance (FCL) necessary to work on classified aspects of a POR. The model established by initiatives such as DARPA's BRIDGES, which specifically sponsors facility clearances for innovative small firms to connect them directly to classified DoD research and development efforts, demonstrates a successful blueprint (Bryant, 2025). Replicating this model ensures that FCL requirements do not become an insurmountable barrier to transition, especially for Path A efforts which are integrated into sensitive programs.

Analytical Framework and Testable Hypotheses

To substantiate the proposed structural transformation of SBIR to SBER, a rigorous, data-driven analytical framework is required. The research design must shift the definition of "success" from R&D completion to measurable acquisition outcomes.

Research Design and Data Plan

The methodological approach requires linking disparate federal data sources to trace the full lifecycle of the SBER investment, from topic initiation to fielded capability.

The foundation of the quantitative analysis lies in joining specific SBER awards with subsequent contract actions. This linkage allows for the accurate measurement of time-to-production and the total monetary value of post-SBER follow-on funding or contract type (SBIR Phase III, Other Transaction Authority, or FAR-based contracts). This data-driven approach is essential, particularly as legislators have expressed a need for agencies to measure and report the transition outcomes of SBIR/STTR funding to quantify its success. Specifically, the analysis must utilize Federal Procurement Data System subcontract flags to identify successful Path A transitions where the small business became a managed component supplier to a prime contractor.

Complementary data must be drawn directly from POR documentation. A detailed review of POR budget exhibits (PPBE artifacts such as R-Docs or P-Docs), System Engineering Plans (SEPs), and Test and Evaluation Master Plans (TEMPs) is necessary to confirm that the SBER technology was truly inserted into the formal acquisition baseline, rather than remaining an isolated test effort.

To isolate the effects of the new governance structure sponsorship as a partnership between the PM, the RO and the SB, a Difference-in-Differences methodology should be employed (Lechner, 2011). This approach would compare the transition outcomes of legacy Phase II awards (the control group) with the new, sponsored SBER awards (the treatment group). Controlling independent variables such as the timeline of mandatory MOSA implementation and the specific service component ensures that differences in transition rates are attributable to the SBER structural changes.

Transformation for the Long-Term: Policy Implementation Roadmap, Risks, and Mitigation

The transition from a research-centric SBIR culture to an acquisition-focused SBER model requires a structured, phased implementation roadmap, accompanied by specific mitigation strategies for anticipated organizational and contractual resistance.

Implementation Timeline for Rapid Concept Integration

The roadmap establishes key milestones to ensure swift institutional adoption and integration into the defense acquisition ecosystem.

Immediate action is required to demonstrate commitment and force cultural change. This phase involves administrative mandates for 90-Day Actions (Administrative and Cultural Shifts):

1. Topic Conversion and Sponsorship: Mandatorily convert a minimum of 20% of the next solicitation cycle’s SBIR topics into the new SBER format, immediately requiring PM and RO mandatory sponsorship for proposal acceptance.
2. Publishing Insertion Calendars: All Portfolio Acquisition Executives must publish mandatory technology insertion calendars for their subordinate programs, establishing a drumbeat of opportunities to align potential SBER solutions with periodic fielding opportunities, both providing clarity and compressing the planning horizon.
3. Infrastructure Refresh: Conduct a rapid refresh of government-stewarded conformance benches and reference architectures to validate the foundation for Path B operations.

12-Month Actions (Pilot and Contracting Reform)

This phase focuses on piloting the dual-path model and formalizing the contractual mechanisms:

1. Pilot Execution: Execute the first set of SBER pilot transitions—roughly half focused on Path A Subcontract Integration (leveraging POR execution) and the other half focused on Path B Incumbent Displacement (leveraging open interfaces).
2. Standardized Contracting: Formalize and mandate the use of standardized Intellectual Property (IP)/data-rights templates for all SBER awards. These templates must clearly define the government's rights to foster competition and ensure successful Path B transitions.
3. Bridge Contracting Utilization: Establish mandatory metrics for the utilization of Middle-Tier of Acquisition (MTA), OTA, and Commercial Solutions Offerings (CSO) authorities to transition SBER Phase II prototypes into Phase III contracts, ensuring robust funding bridges are active across the components.

24-Month Actions (Scaling and Institutionalization)

The final phase involves institutionalizing SBER as the default mechanism for capability introduction:

1. Budget Exhibit Alignment: Mandate the systematic alignment of all R-Docs and P-Docs (budget exhibits) to reflect SBER as a key, mandatory source of POR capability increments, ensuring funding stability for 6.4/6.5 RDT&E and initial procurement planning.
2. Policy Scale-up: Scale the SBER policy across all participating components of the Department of the Navy (DoN), transitioning the entire SBIR budget allocation above a minimal R&D baseline toward the $10 million–$20 million SBER model.

Risk & Mitigation Strategy

Implementing structural change of this magnitude involves significant organizational resistance and logistical risk. These must be proactively mitigated.

Risk 1: Prime Contractor Resistance

Prime contractors (traditional or new entrants), who have traditionally benefitted from closed architectures and proprietary supply chains, will make partial commitment or outright resist the introduction of government-mandated, open-standard components, particularly those enabling Path B Incumbent Displacement.

- *Mitigation:* Resistance is mitigated by tying integrated financial incentives to the prime contractor's adoption of open interfaces and successful management of SBER subcontractors (Path A). For Path B, strict enforcement of MOSA mandates (Title 10 U.S.C. 4401(b)) ensures that resistance to open standards results in competitive disadvantage, forcing compliance with the government-stewarded architecture.

Risk 2: Small Business Capacity and Classification Barriers

Small businesses may lack the internal Systems Engineering (SE) rigor, testing facilities, or the requisite security clearances (Facility Clearance Level - FCL) necessary to manage a large, acquisition-grade contract embedded in a POR.

- *Mitigation:* The SBER award funding itself must include allocations for technical assistance programs, providing embedded SE support and consulting to the small firm. Furthermore, the government must adopt and scale successful models like DARPA's BRIDGES initiative, which provides FCL sponsorship and access to classified work areas, systematically eliminating security barriers for innovative small companies.

Risk 3: Budget Stability for Infrastructure

Funding for foundational infrastructure—reference architectures, conformance benches, and government-stewarded open standards teams—is often vulnerable to discretionary cuts during budget cycles, undermining the viability of Path B.

- *Mitigation:* Funding for these open-architecture infrastructure components must be mandated and categorized not as R&D or RDT&E 6.1/6.2, but as essential acquisition portfolio infrastructure within the 6.4/6.5 or a dedicated, protected budget line item. This classification ensures the tools necessary for managing competition and technical refresh are protected from annual budgetary volatility.

Conclusion and Recommendation Summary: Adding SBER Into Acquisition Transformation

The analysis demonstrates that the current SBIR structure is an adequate seed-fund mechanism for fostering early-stage research but is fundamentally incapable of fulfilling its national strategic goals and legislative mandate to harness the American ingenuity machine in small businesses and transition technology into warfighting capability at scale. The pervasive "SBIR mill" phenomenon is not a failure of small businesses, but a structural incentive failure caused by the misalignment of topics, funding levels, and governance with the complex demands of defense acquisition.

Current legislative reforms may improve transition for a subset of projects by adding higher-end funding pathways and enhanced Phase III authorities. However, they largely overlay the existing SBIR structure rather than redesigning it.

Transition must be embedded in program structure—through acquisition alignment, open modular interface design, funding continuity, and competition mechanisms, so that successful prototypes have a planned path into PORs.

The proposed SBER model provides for structural reorientation by shifting SBIR from early-stage R&D toward engineering, qualification, and integration aligned to acquisition timelines.

Crucially, SBER introduces systemic accountability through its dual transition mandate: Path A (Subcontract Integration) utilizes high-level sponsorship and POR execution to guarantee insertion, while Path B (Incumbent Displacement) enforces competition through government-stewarded Modular Open Systems Approach (MOSA) standards and funded conformance benches.

By implementing the SBER framework, the DoD transforms the small business technology ecosystem from an exploratory feeder of isolated prototypes into a dependable, competitive, and acquisition-ready source of fielded military capability. This structural change is necessary to maintain technological superiority and military readiness by replacing bureaucratic stagnation with predictable, engineering-driven technology insertion.

References

Borda, A. (2025, June 27). Small business programs play critical defense industrial role. *National Defense*. https://www.nationaldefensemagazine.org/articles/2025/6/27/ndia-policy-points-small-business-programs-play-critical-defense-industry-role

Bresler, A. (2023, May). *Assessing the effectiveness of defense-sponsored innovation programs as a means of accelerating the adoption of innovation forcewide*. Acquisition Research Program, Naval Postgraduate School. https://dair.nps.edu/handle/123456789/4868

Bryant. (2025). *BRIDGES initiative: Connecting innovators to classified R&D*. Defense Advanced Research Projects Agency.

Department of Defense. (2024). *Engineering of defense systems guidebook*.

Department of Defense. (2025a). *Acquisition transformation strategy: Rebuilding the arsenal of freedom*. https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/ACQUISITION-TRANSFORMATION-STRATEGY.PDF

Department of Defense. (2025b). *DoD modular open systems approach (MOSA) implementation guidebook*.

Ernst, J. (2026, March 3). *Senate passes Ernst updates to put small businesses first in America's innovation program*. https://www.ernst.senate.gov/news/press-releases/senate-passes-ernst-updates-to-put-small-businesses-first-in-americas-innovation-program

ForwardEdge. (n.d.). *Other transaction authorities (OTAs), CSOs, and SBIR Phase III: What's the difference?*

Gallo, M. (2025). *Small business research programs: Selected issues for reauthorization* (CRS Report No. R48629). Congressional Research Service.

Government Accountability Office. (2024a). *SBIR and STTR: New standards may have minimal effects on multiple award winners* (GAO-24-106398).

Government Accountability Office. (2024b). *Small business research programs: Increased performance standards likely affect few businesses receiving multiple awards* (GAO-24-106398).

Government Accountability Office. (2025). *Small business research programs: Clearer guidance could improve data on topic types and award competitiveness* (GAO-25-107942).

Guertin, N., & Kaliz, M. (2005). *Submarine combat control system development with a focus on human systems integration*.

Lechner, M. (2011). *The estimation of causal effects by difference-in-difference methods*. Department of Economics, University of St. Gallen.

McNamara, R. (2025, August 27). Commentary on SBIR reauthorization: Experienced small businesses play vital role in major acquisition programs. *National Defense*. https://www.nationaldefensemagazine.org/articles/2025/8/27/2-commentary-on-sbir-reauthorization-experienced-small-businesses-play-vital-role-in-major

National Academies Press. (2000). *Phase III transition: Critiques and proposals*.

National Library of Medicine. (2024, July). *Advances in difference-in-differences methods for policy evaluation research*. https://pmc.ncbi.nlm.nih.gov/articles/PMC11305929/

Naval Submarine League. (2012, January). *Team submarine—Providing undersea assets to the warfighter*. https://archive.navalsubleague.org/2012/team-submarine-providing-undersea-assets-to-the-warfighter

Office of the Under Secretary of Defense for Research and Engineering. (2025a). *Modular open systems approach (MOSA) implementation guidebook*.

Office of the Under Secretary of Defense for Research and Engineering. (2025b, February). *Implementing a modular open systems approach in Department of Defense programs*. https://www.cto.mil/wp-content/uploads/2025/03/MOSA-Implementation-Guidebook-27Feb2025-Cleared.pdf

Rothzeid, D. (2025, September 8). *SBIR mills are draining America's innovation fund*. DefenseScoop. https://defensescoop.com/2025/09/08/sbir-mills-are-draining-americas-innovation-fund/

Small Business Innovation Development Act of 1982, S. 881 (1982).

Small Business Innovation and Economic Security Act, S. 3971 (2026).

Stangler, D. (2022). *Should SBIR reauthorization include limits on multiple award winners?* Bipartisan Policy Center.

U.S. House of Representatives Committee on Small Business. (2024). *Chairman Williams introduces legislation to extend the SBIR/STTR programs*.

VanRoo, B. (2022). *Are a few dozen SBIR mills sucking the air out of small business innovation?* https://benvanroo.substack.com/p/are-a-few-dozen-sbir-mills-sucking

Venture Studios to Scale Warfighting Concepts to Capabilities

Dr. Peter Khooshabeh—studied psychological and brain science at the University of California. He authored more than 60 peer-reviewed publications in areas of cognitive science and Human Machine Integration (HMI). His research spans government labs at Department of Energy (DOE), NASA, Air Force Research Lab, and industry, e.g., IBM Research. He was competitively selected as a Defense Ventures Program Fellow to strengthen private sector relationships with Department of War (DoW) innovators and identify ways to partner on technical investments towards capability developments. His current role is Regional Lead of DEVCOM ARL West to integrate science from across trusted academic partners, various industry, and early-stage venture backed companies. [Peter.Khooshabehadeh2.civ@army.mil]

Dr. Michael Bakas—is with the DEVCOM Army Research Office (ARO) was previously program manager for ARO's Synthesis and Processing program, which funds Army scientific research into structural materials processing. He obtained his PhD in 2006 in Ceramic and Materials Science Engineering from Rutgers University by studying the influence of processing-related flaws on the ballistic performance of silicon carbide armor tiles. He continued studying armor materials and their processing at the Idaho National Laboratory, working with transparent ceramics, armor ceramics, and steels. As such, he has experience in both performing and supporting scientific research.

Mr. Michael D'Onofrio—is a Technology Transfer Specialist at the DEVCOM ARL. He specializes in transitioning the Army's cutting-edge intellectual property to the commercial sector through mechanisms such as Cooperative Research and Development Agreements (CRADAs), patent licensing, and novel programs such as the Mission Transition Series. With more than three decades of experience in technology transfer, he plays a critical role in bridging the gap between laboratory innovation and real-world commercial and military capabilities, ensuring the Army's research drives the development of new products. He holds an MBA and a BS in Mechanical Engineering from the University of Maryland.

Abstract

Venture studios are rapidly emerging as efficient entities for creating companies, developing innovative business concepts, evaluating market traction, providing seed funding from associated venture capital (VC) firms, and accelerating growth through strategic partnerships. This paper explores how the venture studio model can uniquely support the scaling of emerging technologies into operational warfighting capabilities. By constructively analyzing various venture studio models and identifying best practices, this research aims to provide actionable insights for the broader National Security science and technology ecosystem. Using a case study methodology, the paper examines successful venture studio exemplars to extract lessons learned and strategies for effective technology transition across the broader government science and technology laboratory ecosystem.

Problem Statement and Background

Transitioning emerging technologies out of laboratories and into the hands of users is a daunting challenge. Doing this in a military context proves that much more difficult. What are proven innovation methodologies that increase the transition probability of successful lab-based emerging technologies? In this report we analyze the role of venture studios in tech transition across the government, academic, and industrial science and engineering nexus.

Venture studios represent a significant evolution in the entrepreneurial ecosystem, shifting the paradigm from an entrepreneur-driven activity to a systematic, organization-level process of venture creation. Patel and Chan (2024) describe this as a move toward an "'assembly line' serialization of venturing in an organization" (p. 285). Also referred to as venture builders, startup factories, or company builders, venture studios are a novel form of startup-supportive organization that internally generates, validates, builds, and funds concurrently multiple new business ideas (Patel & Chan, 2024). Unlike traditional support models, the studio itself is the central actor, leveraging in-house resources and a dedicated team to “create, develop, fund, and spin-off internal startups from scratch” (Patel & Chan, 2024, p. 286).

While often grouped with incubators and accelerators, the venture studio model is operationally distinct. Patel and Chan (2024) emphasize that compared to accelerators, which support cohorts of externally formed startups, venture studios take a “direct, hands-on approach” to building a venture within the boundaries of the studio itself (p. 284). A key differentiator is the studio’s deep involvement in the earliest stages of ideation and team formation. Furthermore, studios typically take a significantly larger equity stake (30%–80%) than accelerators (5%–10%), a reflection of their intensive, in-house process of de-risking the venture from its inception (Patel & Chan, 2024). This model institutionalizes the venturing process, providing a stark contrast to market-based support relationships.

From a theoretical perspective, the venture studio model can be understood through the lenses of transaction cost economics and the resource-based view of the firm. Patel and Chan (2024) argue that by internalizing the venture development process, studios significantly lower the transaction costs associated with coordinating between disparate actors like entrepreneurs, investors, and developers. This internalization allows for better alignment of incentives, improved monitoring, and more efficient resource allocation (Patel & Chan, 2024). From a resource-based view, studios build a unique, fungible bundle of resources—including expert human capital, repeatable processes, and established networks—that can be leveraged and shared across a portfolio of ventures. This creates economies of scope and provides a scalable, systematic framework for innovation that is difficult for individual entrepreneurs to replicate (Patel & Chan, 2024).

Dual-Perspective Analysis of the Current Tech Transfer Model

There is a strong potential for the venture studio model to be a novel pathway for transitioning intellectual property (IP) from Department of War (DoW) laboratories broadly, to include academically funded work and industry partnerships, to broader commercialization and markets. Hence, venture studios offer a potential solution to the persistent “valley of death” in defense innovation. However, the efficacy of this model must be evaluated from the dual perspectives of the two primary stakeholders: the technical inventor (the source of the IP) and the venture (the commercialization and related investment vehicle). This analysis examines the distinct advantages and disadvantages for each party, reconciling the theoretical benefits of the studio model with the practical realities of the DoW technology transfer ecosystem.

The Inventor’s Perspective: A Non-Traditional Path to Impact

For the technical inventor, engaging with a venture studio represents a non-traditional but potentially high-impact pathway for technology transition. The primary benefit is the opportunity to gain direct, real-world experience in the commercialization process as some sort of inventor or entrepreneur in residence. Inventors are potentially exposed to industry partners (e.g., as future customers or acquirers of the resulting tech transfer entity or product, market validation techniques, and the inherent risks of building a business) providing them with an invaluable perspective that is often absent in a laboratory environment. This engagement expands their professional network and provides a firsthand understanding of the challenges small companies face when attempting to build a business around government funded IP.

However, this pathway is not without its challenges for the inventor. A significant con is the potential for a loss of control, as the venture studio may pivot the technology in a direction the inventor had not originally intended but that could prove more fruitful in gaining broader market traction. Furthermore, structural and bureaucratic hurdles present significant friction. Lab funding is typically structured for specific research use, and there is often no formal mechanism or dedicated funding to support the inventor’s continued engagement with the venture beyond the initial studio program. While entrepreneurial leave is an option, it often requires the inventor to forgo their government/academic salary, creating a significant personal financial risk that can deter participation. This highlights a critical gap between the desire for tech transfer and the practical means to support the human capital required to make it successful.

The Venture's Perspective: Access and Obstacles

From the venture's standpoint, the studio model provides privileged access to a curated portfolio of government-funded IP and, crucially, direct access to the inventor during the program. This proximity to the subject matter expert is a significant de-risking factor. The association with a federally funded program also lends credibility and can be a valuable asset when engaging with early-stage stakeholders and potential investors. The model provides a clear pathway to formalize the relationship through established tech transfer mechanisms, such as CRADA or a licensing agreement, after the initial studio phase.

Despite these advantages, the venture faces considerable obstacles inherent to the government tech transfer process. The timeline for executing a CRADA can be relatively lengthy and often requires an "in-kind" contribution from the startup, a significant demand on a resource-constrained entity in its earlier stages. Hence, a more apt CRADA partner may be a venture studio or even a fund that is considering deep/defense technology in their horizon. A more fundamental challenge is the age and availability of the IP. By the time a patent is officially issued by the U.S. Patent and Trademark Office (PTO), the underlying technology may be 3–4 years old, risking obsolescence in a fast-moving market. Compounding this, the original inventor may have moved to a different project, left the organization, or become otherwise unavailable, creating a "cold start" problem for the venture. Finally, there is a frequent misalignment on IP rights. Most ventures seek an exclusive license to justify their investment, whereas government labs are often reluctant to provide exclusivity, preferring to grant a more limited "field of use" license. This fundamental conflict over exclusivity remains a major friction point in the DoW tech transfer landscape.

Method

We used a structured interview protocol with top venture studios that intersect with the defense technology domain. The Army Human Research Protection Office (AHRPO) reviewed the interview protocol for this study.

Results

Best Practices and Lessons Learned from Share Ventures: A Case Study

An interview with Mr. Hamet Watt, founder of venture studio Share Ventures (www.share.vc), provides critical insights into the operational methodologies that drive successful deep tech commercialization (LA Times Studio Staff, 2025). The studio's approach is built on a foundation of data-driven sourcing, a focus on market creation, and a sophisticated understanding of talent and team dynamics. The following analysis synthesizes the key best practices and lessons learned from this engagement.

Technology Sourcing and Evaluation: From Relationships to Data-Driven Discovery

A primary lesson from Share Ventures is the evolution of their technology sourcing strategy. While traditionally reliant on personal networks and relationships, the studio has increasingly adopted a technology-first approach. They employ a proprietary AI-powered system that continuously scans more than 35 different data sources, including the latest scientific papers, to identify emerging research in their focus areas, such as human performance. This is conceptually similar to the approach of Deep Science Ventures with their Elman tool as reviewed by Khooshabeh et al. (2025).

However, the evaluation of this technology goes beyond traditional metrics like Total Addressable Market (TAM). The studio's core evaluation criterion is the search for a technological "unlock"—a capability that can democratize a product or service, thereby creating a market far larger than what currently exists. As Mr. Watt explains, a myopic focus on the existing taxi industry's TAM would have missed the half-trillion-dollar opportunity unlocked by ride-sharing. This forward-looking, market-creation mindset is a critical departure from conventional market analysis.

Furthermore, the studio explicitly prices in the "degree of difficulty" for deep tech. If the technical complexity, cost, or development timeline is too high, the opportunity receives a lower score, even if scientifically promising. This demonstrates a disciplined, risk-adjusted approach to prioritizing ventures versus simply extending fundamental science that may be less likely to transition as a viable or scalable technological capability.

Talent and Team Formation: The "Lightning Rod" Principle

Share Ventures' strategy for recruiting technical experts and entrepreneurs is not based on broad outreach, but on a focused "lightning rod" principle. The key is to become exceptionally good at defining and articulating a specific, compelling problem.

Mr. Watt emphasizes that when a studio can clearly visualize and communicate a problem it is passionate about solving, it acts as a magnet for top-tier researchers and entrepreneurs who are already obsessed with that same problem. This shared passion creates an "instant bond" and a foundation of trust that is far more powerful than a simple transactional relationship. This approach also acknowledges a critical challenge in deep tech transition: the frequent disconnect between research expertise and product expertise. By starting with a clearly defined problem, the studio can more effectively build a team that combines the necessary scientific depth with the commercialization and product skills that a venture studio brings to the table and are required to build a viable business.

Venture Management: The Codified Playbook and AI-Powered Execution

For ventures with long development times and high technical risk, Share Ventures relies on a highly structured and codified "playbook." Every venture is broken down into granular workstreams—such as Product, Demand (marketing/sales), Team (talent/culture), and Operations (finance/legal)—all work that a very early startup might otherwise neglect but for which a venture studio provides fractional shared human resources to support early ventures. While seemingly mundane, these workstreams of Product (design etc.), Demand generation with (go-to) market/sales, Team formation with the appropriate talent that have the best cultural fit as well as legal/financial Operations prove crucial to whether a compelling deep/defense tech innovation from a lab may scale to broader capability.

A key venture studio management principle is to prioritize the "tall pole in the tent"—the hardest, most value-driving problem—first. This ensures that resources are focused on derisking the most critical aspect of the venture from the outset. This entire process is supported by a central system of approximately 150 AI models that assist in task breakdown, execution, and resource allocation. This demonstrates a systematic, repeatable, and data-driven approach to venture building, treating the studio not as a collection of ad-hoc projects, but as a true "startup factory."

Ecosystem Development: Beyond Meetings to Shared Experiences

The studio's approach to building partnerships with industry and academia moves beyond traditional networking. Two key strategies stand out:

Engage Acquirers Early: Mr. Watt highlights the importance of building relationships with the corporate development community at large technology companies (e.g., Google, Amazon) from the very beginning. These are the ultimate buyers of successful startups, and their insight into the problems they are trying to solve is invaluable market intelligence.

Create Shared Experiences: Instead of standard meetings, the Share Ventures studio focuses on creating unique, themed "experiences" like curated dinners or hackathons. These events are designed to move participants out of their "default" professional personas and reveal their true passions and capabilities. A hackathon, for example, allows the studio to see potential talent "in action," providing a much deeper assessment than a traditional interview. The Defense Hackathon Act of 2024 was passed by the U.S. Congress and could provide opportunities to synergize DoW authorities to spur tech transition with private capital partners. It requires each of the services to put on four hackathons every year and these technical venues could be viable windows for both venture capital and the government to perform technical due diligence on promising teams or capability developers.

Key Lesson for DoW Tech Transfer: The Primacy of the Problem Statement

When asked what one thing could be changed to better support deep tech transition from government labs, Mr. Watt's answer was unequivocal: a clearer, more accessible feed of the problems the Army is trying to solve and the related demand signal associated with solving those problems. He argued that innovators and entrepreneurs are drawn to solving hard problems. When the problem is ambiguous or difficult to find, talented people waste immense energy on defining a problem on their own rather than problem solving what Warfighters or other subject matter experts could more efficiently synthesize and articulate. A simple, well-articulated problem statement (e.g., "We need to solve back pain for soldiers, which costs the Army billions") acts as the most powerful "lightning rod" for attracting the nation's top talent and private capital.

Best Practices and Lessons Learned from Red Cell Partners

An interview with Red Cell Partners (www.redcellpartners.com), a venture fund with an associated venture studio, offers a distinct perspective on the venture studio framework, particularly as it applies to complex, regulated markets such as healthcare and national security (Knapp, 2023). Red Cell Partners' methodology is defined by its strategic focus on "hard" markets, its "platform" approach to venture building, and its flexible, multi-pronged strategy for company formation.

Strategic Focus: Targeting "Bureaucratic Markets"

Red Cell Partners deliberately targets highly regulated and bureaucratic markets; specifically healthcare, cybersecurity, and national security. Their venture studio's core thesis is that while these markets are difficult to penetrate due to red tape and the need for deep institutional knowledge, they also hold the greatest opportunity for disruption. This focus informs their entire model, which is designed to provide startups with the necessary network and operational support to navigate these complex ecosystems effectively.

The "Platform" Model: A Shared Services Foundation

A key differentiator for Red Cell Partners is its "platform" approach. The studio maintains a large core team of 70–80 professionals providing a comprehensive suite of shared services, including engineering, human resources, talent acquisition, accounting, finance, and marketing. When a new company is formed or brought into the venture studio, this platform acts as its "business foundation." The Red Cell Partners venture studio allocates team members to perform these essential functions for the startup. This model is designed to solve a critical early-stage problem: it allows the startup's core team to focus exclusively on high-value activities like customer discovery, go-to-market strategy, and technology development, while the venture studio's platform handles the "annoying" but necessary operational overhead. This approach is intended to "hyperscale their growth" by removing common administrative and bureaucratic hurdles.

Red Cell Partners does not adhere to a single method for company formation. Instead, it employs a flexible, opportunistic strategy that includes three distinct pathways. First, they focus on company building from within. Their venture studio identifies a need and builds the solution entirely in-house with its own product and engineering teams. The company DEFCON AI, focused on leveraging AI for contested logistics, is a prime example of this model.

The second method for their company formation is to acquire/hire to accelerate solving a problem within the venture studio. Red Cell Partners identifies a very early-stage company (e.g., a two-person team with a prototype) and brings them in-house through an "acquihire." The venture studio then puts its full platform team behind the existing company to reshape the product and accelerate its growth.

A third approach to company formation is technology licensing. Their venture studio identifies a market need, scouts for relevant technology in universities and labs, and then builds a company around licensed IP. Their portfolio company DeployX was formed this way, licensing a novel, small form-factor antenna technology from the University of Toledo to address a need in the electromagnetic spectrum. The research behind the university IP has prior DoW funding, e.g., from the Office of Naval Research (ONR; Van Winkle, 2021). This multi-pronged approach to company formation allows Red Cell Partners the flexibility to be open to any and all opportunities, starting with the market need first and then determining the best path to build a company around it.

Red Cell Partners utilizes a structured, multi-phase process to de-risk its ventures through a phased incubation and "Slow Off-Boarding" process. In first phase, a team of analysts spends 3–4 months conducting deep-dive research into broad opportunity areas (e.g., electronic warfare, undersea capabilities) to identify specific, interesting problems. The team next spends another 3–4 months during a discovery phase once they've identified and selected a problem that requires solving by the venture studio. This "Discovery" sprint focuses heavily on customer discovery and results in a comprehensive 12-month plan for a potential company formation, including competitive analysis, a tech roadmap, and a go-to-market strategy. Next, following a "go" decision from the investment committee, Red Cell Partner puts forth its own $1 million–$2 million of funding to provide the newly conceived company approximately eight months of runway. After the startup executes on its plan and demonstrates progress, Red Cell Partners invests a second tranche of $2 million–$3 million. Lastly, at least during the new companies time within the venture studio, it seems to graduate via external validation. Hence, the goal is for the startup to reach a critical inflection point where it has enough traction, e.g., a Minimum Viable Product (MVP) or early proofs of concept, to raise a seed round from an external lead investor. This external validation is the key graduation trigger. The process is described not as a hard stop, but as a "slow off-boarding," where the company gradually hires its own team and reduces its reliance on the Red Cell Partners platform as it scales.

Best Practices and Lessons Learned from Roadrunner Venture Studios: A Case Study

An interview with Steve Weinstein—a former Naval Research Lab scientist, serial entrepreneur, and current Chairman of the Roadrunner Venture Studio—provides a candid and critical analysis of the challenges and opportunities in transitioning deep technology from government labs. The operational model of Roadrunner (www.roadrunnerventurestudios.com), a state-funded venture studio in New Mexico, is built on a set of principles that directly address the common failure points of traditional government technology transfer (Mathews, 2024).

A foundational lesson from the Roadrunner model is its deliberate rejection of an "IP-first" sourcing strategy that represents its core philosophy of first identifying challenging problems. Weinstein was unequivocal that simply browsing a laboratory's inventory of available patents "doesn't work well for us" (S. Weinstein, personal communication, March 6, 2026). He argues that this approach, exemplified by initiatives like the "patent holiday," is misaligned with how venture creation actually works, because "patents are not products and people invest in businesses, not patents." Even so, recent reporting suggests that in just the few months since the announcement of the patent holiday, there has been strong interest and negotiated licensing terms (Luckenbaugh, 2026).

Instead of pursuing IP first, Roadrunner employs a thesis-driven, problem-first approach. The studio begins by identifying a broad area of interest (e.g., geothermal energy) and then narrows its focus to a specific, high-potential problem. Only after the problem is clearly defined does their venture studio begin the process of finding the right people and technology to solve it. This methodology inverts the traditional government tech transfer model, which often pushes technology out of the lab in search of a problem.

Central to Roadrunner's strategy is its use of the Entrepreneur-in-Residence (EIR) model making it truly founder-centric. Rather than starting with a technology, the studio's first step is to "double down on finding a person who can become an entrepreneur in our studio" (S. Weinstein, personal communication, March 6, 2026). As part of their EIR development process, an individual with deep domain expertise and entrepreneurial drive is brought into the venture studio on a stipend for a period of approximately six months. This engagement is structured as a "cancelable purchase order," providing flexibility if the fit is not right. The EIR's primary task is to validate the venture studio's thesis about a previously broad area of interest and then "go shopping" for the right technology at universities or national labs. The selection criteria of an EIR are based on trust and coachability. Weinstein emphasizes the importance of finding someone who will "die on the hill with our money" and who is willing to take feedback. This founder-centric approach ensures that a passionate, dedicated leader is driving the venture from day one, a critical factor often missing when a technology is simply licensed out of a lab without the inventor necessarily coming with the IP license. Hence, a major critique of the government's approach is its focus on the patent as the primary unit of transfer. Weinstein argues that for a startup, a patent alone is insufficient and often unattractive. A successful transition requires a complete "package" that includes not just the IP, but also access to the scientists who created it and the specialized test equipment used to develop it. Without this holistic package, the Technology Readiness Level (TRL) is often too low and the friction too high for a startup to succeed.

The most significant structural barrier identified is the "clock speed" mismatch between startups and government labs. Weinstein notes that a startup operates on a timeline of months, while the process to license technology from a federal lab can take about a year or more (National Institute of Standards and Technology et al., 2023). For a startup, "a year is an infinite time" (S. Weinstein, personal communication, March 6, 2026). This lengthy and unpredictable transaction time may make it less attractive for an entrepreneur to build their critical path around government IP. As a result, startups will often choose to work independently versus licensing IP from government labs, even if it means working around existing government patents, because the risk of delay is greater than the risk of potential infringement.

While there are temporal risks associated with a nascent or yet to be started new company licensing IP from government labs, Weinstein indicates the significant value in the DoW's science and technology ecosystem. He argues that the most valuable signal for the venture community is not an issued patent, but the initial funding of a research project. By the time a patent is issued, VCs are "probably two years too late." The fact that agencies like ARO, ONR, and AFOSR have used its internal experts to vet and fund academic talent is a powerful, early endorsement of the research's potential. A simple, searchable, and regularly updated newsletter or database highlighting these funded projects would be an invaluable resource for the VC community, providing a much earlier and more effective signal for identifying promising technologies than any list of available patents, i.e., a leading indicator of sorts.

Discussion

The venture studio model, while gaining prominence, is not a monolith. An analysis of the operational strategies of three distinct studios—Share Ventures, Red Cell, and Roadrunner Venture Studios—reveals a spectrum of approaches to sourcing, building, and scaling deep technology ventures. While all three embody the core "startup factory" concept (Patel & Chan, 2024), their differing philosophies on opportunity identification, team formation, and operational support provide a rich tapestry of best practices. Contrasting these models offers a nuanced playbook for how the DoW science and technology ecosystem can better leverage this framework for technology transfer.

Sourcing and Opportunity Identification: Problem vs. Market vs. Founder

The initial stage of venture creation highlights a key philosophical divergence between the studios. Share Ventures employs a problem-first "lightning rod" principle. The studio invests heavily in defining and articulating a compelling, often underserved, problem (e.g., using oral health as a predictor for systemic health). This clear problem statement then acts as a magnet to attract passionate experts and entrepreneurs. This process is increasingly augmented by a data-driven engine that scans more than 35 sources to identify emerging research trends that align with their problem focus.

In contrast, Red Cell adopts a market-first approach. The studio strategically targets "bureaucratic markets" like national security and healthcare, where high barriers to entry create opportunities for disruption. Their process begins with a top-down analysis of demand signals, such as tracking capital flows in congressional bills, to forecast where funding will be allocated. This market analysis dictates the specific problems they choose to solve.

Roadrunner Venture Studios presents a third, distinct model: a founder-first approach. As articulated by Chairman Steve Weinstein, the process begins not with a problem or a technology, but by "doubling down on finding a person who can become an entrepreneur in our studio" (S. Weinstein, personal communication, March 6, 2026). Roadrunner Venture Studios then applies their platform of foundational and shared startup formation platform to empower the EIR to validate a thesis and "go shopping" for the right technology to solve a problem within that thesis. This model inverts the traditional tech transfer process, prioritizing the human capital of the founder above all else.

Venture Creation and Operational Support: The Platform Spectrum

Once a venture is initiated, the venture studios provide intensive support, but the nature of this support exists on a spectrum. Red Cell represents the most comprehensive platform model. With a large in-house team of 70–80 professionals, the studio provides a full suite of shared services, including engineering, finance, legal, and HR, that act as the startup's business foundation. This allows the startup's core team to focus exclusively on go-to-market strategy and customer discovery. The goal is to hyperscale growth by removing administrative friction, with the startup gradually off-boarding from the platform as it raises external capital and hires its own team.

Share Ventures employs a more process-oriented platform, utilizing a codified playbook and a suite of AI tools to manage workstreams and prioritize the tall pole in the tent which represents the venture's hardest problem. While still providing deep operational support, the emphasis is on a repeatable, data-driven management system rather than a large, dedicated service organization. Roadrunner's model, being founder-centric, empowers the EIR with initial capital and strategic guidance, but appears less focused on providing a large, centralized service platform, instead prioritizing the EIR's autonomy in building their own team and processes.

Lessons for the DoW S&T and Acquisition Ecosystems

Contrasting these three successful but different models provides several critical lessons for the broader DoW tech transfer community and acquisition professional communities. First is the primacy of the problem statement. The most powerful tool for attracting external talent and capital is a clearly articulated set of problems. The DoW's current method of publishing available IP is a promising start to spur tech transition but could provide more useful by aligning the inventors who could articulate the warfighting problems they were solving in their work that eventually led to the formation of the IP. To effectively engage with the venture ecosystem, the DoW must create a "clear feed" of its most pressing operational needs, acting as a "lightning rod" for innovation.

A second lesson learned is the Value of a platform of foundational services for navigating bureaucracy. This platform model offers a compelling solution to the immense bureaucratic hurdles that startups face when trying to work with the DoW and also scaling their companies. A government-sponsored or partnered entity that could provide centralized, shared services for contracting, security compliance, e.g., managing access to Controlled-Unclassified Information (CUI), etc., and legal review could dramatically accelerate the ability of small, innovative companies to transition their technology to the DoW.

A third lesson is that prioritizing the value of venture studios being able to license or access package of talent, government testing equipment, etc., over licensing just a patent. The government's focus on licensing individual patents is fundamentally misaligned with the needs of a startup. As Weinstein noted, access to government scientists and access to the test equipment may help better align with a startups clock speed. Where government licensing can take on average just under a year (National Institute of Standards and Technology et al., 2023) while a startup operates on a timeline of months, so this temporal mismatch could be a critical barrier. In addition to IP licensing, other partnership vehicles such as CRADAs could more quickly align with startups' clock speeds. This suggests that the DoW could explore founder-centric models, such as sponsoring its own EIRs, and radically streamline its tech transfer processes to operate at the speed of business, not bureaucracy.

In conclusion, while no single venture studio model is a panacea, a hybrid approach that combines a clear, problem-first sourcing strategy, a flexible set of venture creation tools, and a platform-based support system offers a powerful framework. By adopting these best practices, the DoW can more effectively bridge the "valley of death" and translate its scientific innovations into tangible capabilities for the Warfighter.

Acknowledgments

The authors acknowledge the managing partners and other staff of the venture studios we had the fortune to interview for this manuscript. We used GenAI.mil, which as of this writing primarily includes Google Gemini, to reconcile text transcripts of the interviews with our own notes during those same meetings with venture studio leaders and practitioners.

References

H.R. 10333, 118th US Congress. (2024). *Defense Hackathon Act of 2024*.

Khooshabeh, P., Rose, A., & Campbell, T. (2025). *A paradigm shift for how DoD funds people to drive innovation through entrepreneurial science* [Presentation]. Acquisition Research Program. https://dair.nps.edu/handle/123456789/5364

Knapp, A. (2023). Red Cell partners raised a $91 million fund to bring AI to healthcare and defense. *Forbes*. https://www.forbes.com/sites/alexknapp/2023/12/21/red-cell-partners-raised-a-91-million-fund-to-bring-more-ai-to-healthcare-and-defense/

LA Times Studio Staff. (2025). Hamet Watt on building businesses that enhance longevity and well-being. *Los Angeles Times*. https://www.latimes.com/b2b/innovators-unplugged/story/hamet-watt-interview-building-share-ventures

Luckenbaugh, J. (2026, April 6). Pentagon finds success with patent holiday program. *National Defense: NDIA's Business and Technology Magazine*. https://www.nationaldefensemagazine.org/articles/2026/4/6/pentagon-finds-success-with-patent-holiday-program

Mathews, J. (2024, September 4). New Mexico's sovereign wealth fund is investing $50M in a bet that scientists will build startups in Albuquerque. *Forbes*. https://fortune.com/2024/09/04/new-mexico-sovereign-wealth-fund-50-million-roadrunner-startups-albuquerque/

National Institute of Standards and Technology, National Oceanic and Atmospheric Administration, National Telecommunications and Information Administration, & Institute for Telecommunication Sciences. (2023). *Annual report on technology transfer: Approach and plans, fiscal year 2022 activities and achievements*. https://www.nist.gov/system/files/documents/2023/09/13/Annual%20Report%20on%20Technology%20Transfer%20-%20Approach%20and%20Plans%2C%20Fiscal%20Year%202022%20Activities%20and%20Achievements%5B84%5D_2.pdf

Patel, P. C., & Chan, C. S. R. (2024). The influence of differences between venture studios on differences in venture outcomes. *Venture Capital*, *26*(3), 283–301. https://doi.org/10.1080/13691066.2023.2185168

Van Winkle, D. (2021). UToledo researcher awarded grants to develop plasma-based RF electronic systems. *UToledo News*. https://news.utoledo.edu/index.php/09_21_2021/utoledo-researcher-awarded-grants-to-develop-plasma-based-rf-electronic-systems

The Policy Test Lab: An Agentic AI-Based Simulation Tool

Douglas J. Buettner, PhD—is the AIRC Deputy Chief Scientist at the Stevens Institute of Technology and is an Adjunct Assistant Professor for the University of Utah's Systems Engineering Program. His doctoral research at the University of Southern California's Viterbi School of Engineering examined issues with DoD software-acquisition strategies, identifying acquisition optimizations. His research is found in Acquisition Games and Software Defects: Dynamics of the Software Developer's Dilemma. Dr. Buettner has BS and MS degrees in Physics, with a minor in Atmospheric Science from Oregon State University, and his PhD from USC is in Astronautical Engineering. [dbuettne@stevens.edu, 201-216-8300]

Raz Saremi, PhD—is an AIRC Research Scientist at the Stevens Institute of Technology. Dr. Saremi brings extensive industry experience in the machine learning, data science, and advanced analytics domain, with leadership roles at Archer Daniels Midland (ADM), Johnson & Johnson, Celgene (BMS), and Merck. At ADM, she led the company's generative AI strategy, designing large-scale market intelligence and forecasting systems using Large Language Models, network science, time series forecasting, and multi-objective optimization to improve strategic decision-making across commercial and supply chain operations. Her PhD is in Software and Systems Engineering from Stevens. [rsaremi@stevens.edu, 201-216-8300]

Ethan Silverstein—is a graduate student at the Stevens Institute of Technology. He is a Research Assistant in the Mobile AI Cybersecurity Computing laboratory, focusing on AI Operations for 5G communications systems, formal verification methods for distributed AI systems, and UAV security protocols with emphasis on scalable AI deployment and monitoring. He spent last summer as a Research Intern at the Air Force Research Laboratory designing multi-agent reinforcement learning policies for distributed command and control problems, outperforming traditional heuristic methods. Ethan's BS is in Computer Science from Stevens. [esilvers@stevens.edu, 201-216-8300]

Zoe Szajnfarber, PhD—is the SERC Chief Scientist at George Washington University and is a Professor of Engineering Management and Systems Engineering, International Affairs, and Computer Science, Faculty Director of the GW Trustworthy AI Initiative. Her research group focuses on the design and development of complex systems in their socio-technical context. Recent work has focused on strategies for leveraging nontraditional design inputs, from open innovation to AI-enabled workflows. Dr. Szajnfarber's MS and PhDs are from the Massachusetts Institute of Technology. [zszajnfa@email.gwu.edu, 202-994-6080]

Abstract

As an emerging technology, Agentic Artificial Intelligence or Large Language Models (LLM)–based/Generative AI agent systems, have been increasingly adopted to enable autonomous reasoning, tool-using, and decision-making systems that transcend the traditional boundaries and use cases of language models. Many of these systems rely on multiple interacting agents, each capable of perception, reasoning, and adaptive behavior. Such agents collaborate in multilayered ecosystems to achieve cognitive and operational objectives. The policy test system introduced in this paper extends the agentic paradigm into the policy simulation domain by designing an LLM-based orchestration agent that autonomously interprets academic documents and translates them into executable implementations of simulation and optimization models.

While this research is still in its infancy, the agentic tool already extracts an analogous model from peer reviewed literature, where the LLM serves as a cognitive controller, parsing unstructured knowledge into executable code. The resulting code provides simulation agents with the underlying dynamics of policies, resource flows, and behavioral adaptation. This work is a step towards a future where generative reasoning agents autonomously analyze, simulate, and optimize complex socio-technical systems in support of informed policy exploration.

Introduction

Policy test simulation has been an active area of research by the Stevens Institute of Technology's Systems Engineering Research Center (SERC) for many years (Rouse, 2014;

Rouse & Bodner, 2012, 2013a, 2013b; Rouse et al., 2019; Kannan et al., 2022). However, traditionally the approach requires upwards of 2 years to complete, starting with interviews of the affected stakeholders, developing a model of the policy, finding and using the data for calibration and testing the model, and then exploring the ramifications of those policies on the affected stakeholders. While the approach should still be considered as the "gold-standard," the immediate need for policy change to affect our acquisition transformation strategy across our war fighting enterprise outweighs the time required to fully develop and test these models. This begs the question, *Is there another way?*

The Policy Test Lab (PTL) that is now under development by the Acquisition Innovation Research Center (AIRC) leverages Agentic Artificial Intelligence (AI) or Large Language Models (LLM)–based/Generative AI agents with autonomous reasoning and tool-using capabilities in support of exploring a decision-maker's policy questions. Even though the PTL is in its infancy, it has the potential (within computational constraints of course), to simulate virtually any complex socioeconomic system. Initially, it will be used to simulate and test Department of War (DoW) acquisition transition strategies (DoW, 2025).

Figure 1 shows the Operational View (OV-1) of Rouse's original Policy Flight Simulator (PFS).[1] Once implemented, the PTL is expected to operate in much the same manner but will use Agentic AI to implement many of Rouse's steps. Furthermore, as models are built, explored, and completed for the policy-makers, they will be added to a central repository for future exploration and refinement, becoming building blocks for more complex models. Hence, an OV-1 of the PTL is quite similar, containing the same essential steps under the hood, but now they are automated using skilled AI agents executing this multistep process. A revised OV-1 with the "Multi-Level Computational Policy Simulator" or (MCPS) at its core is shown in Figure 2.

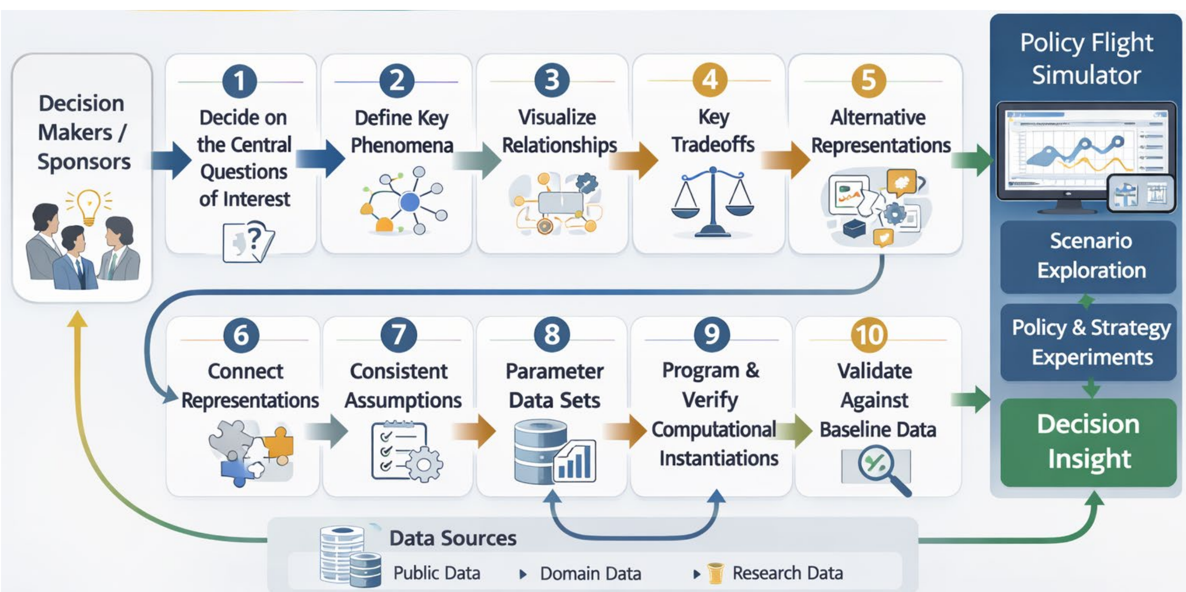

Figure 15. Operational View (OV-1): Rouse's Policy Flight Simulator

[1] Appendix A directly quotes Rouse's (2014) 10-Step Methodology and is provided for reference.

Figure 16. Operational View (OV-1): The Agentic AI-based Policy Test Lab

The remainder of this paper discusses the MCPS design and architecture, and its current implementation status.

The Multilevel Computational Policy Simulator (MCPS) Architecture

Internal to the PTL's MCPS, the task flow starts with a policy question and uses a chat agent to create a scenario that is used to help identify existing literature in the policy question's area, focusing on papers containing models frequently used in policy studies, namely system dynamics, discrete event, and agent-based or hybrid models. The scenario contains key words and synonyms for a detailed literature search, and identifies a list of key stakeholders and suggested names for stakeholder agents in the policy area, associated topics, and model types to look for with potential sources of data. The scenario is provided to the PTL's Orchestrator Agent, which then builds a list of queries for a literature search. The user also has the option of expanding from a generic OpenAI "websearch," to include in each query a targeted organization's name, such as RAND, MITRE, DTIC, SERC, or AIRC, among others.[2] The general concept is to first identify what has been already studied. We decided to start with this very generic approach to allow the tool to use models that have already been built, if they exist, and if nothing is found, then we have the basic information from which to start building our own models from scratch.

Figure 3 is the process view sequence diagram from our initial phase of development.

2 It is important to note that when including organizations, the literature search can find potentially hundreds of articles, which can add up quickly even though we have defaulted to using cheaper models.

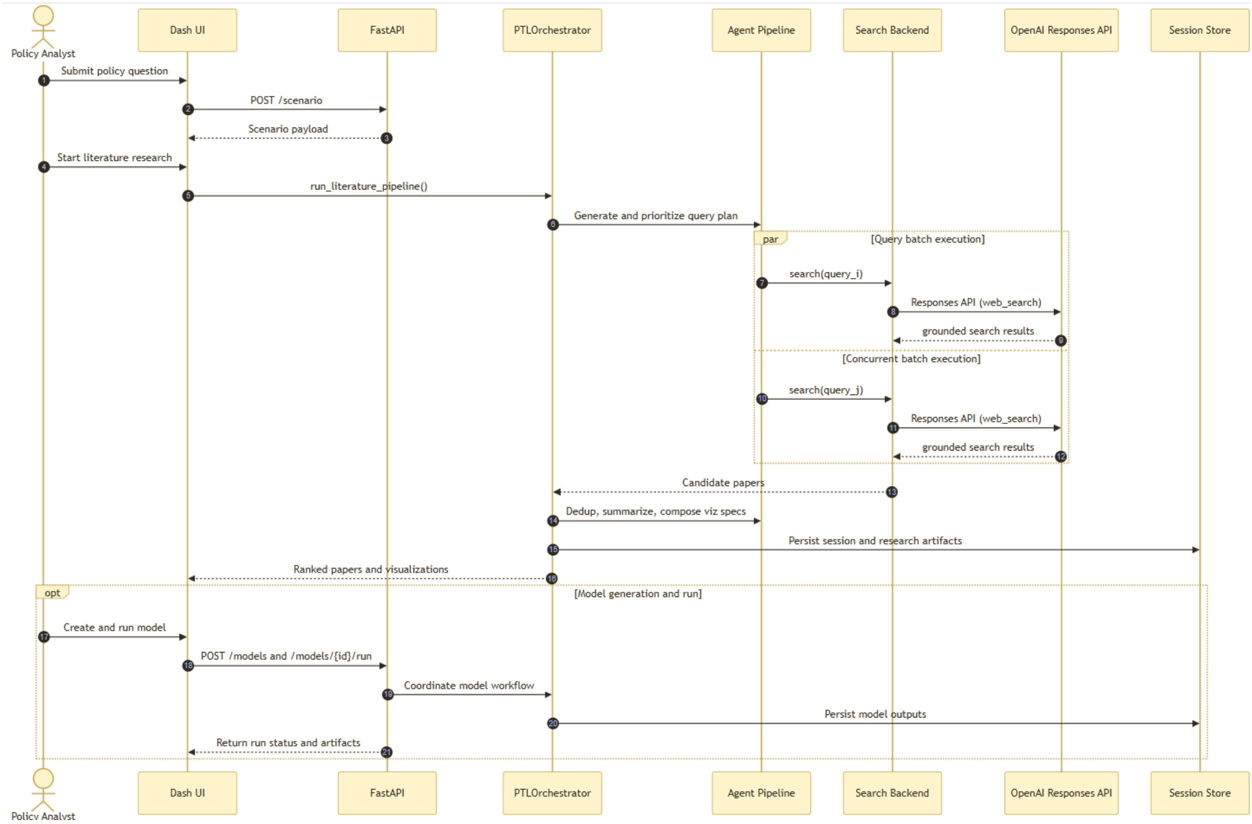

Figure 17. Process (View) Flow for this Version of the Policy Test Lab's MCPS

Select 4 + 1 architectural views are provided in Appendix B.

The PTL's Simulator Generator

If one or more promising papers are identified, the policy analyst can then send that paper to the Simulator Generator suite of agents to carefully analyze the paper's contents for existing models, which includes coding agents that generate Python code for the models found within. To contain costs while the PTL is under development and being debugged, we have restricted these agents from analyzing papers with more than 50 pages of text, and then we have focused on building very simplistic representations of potentially complex models.[3] The Simulator Generator's multi-document upload interface is provided in Figure 4, with Figures ***Figure 20***, ***Figure 21***, and ***Figure 22*** showing extracted System Dynamics Model (SDM) simulation, Agent-Based Model (ABM), and Discrete Event Simulation (DES) outputs respectively.

[3] We expect to include more agents in the future to accurately reproduce a paper's model.

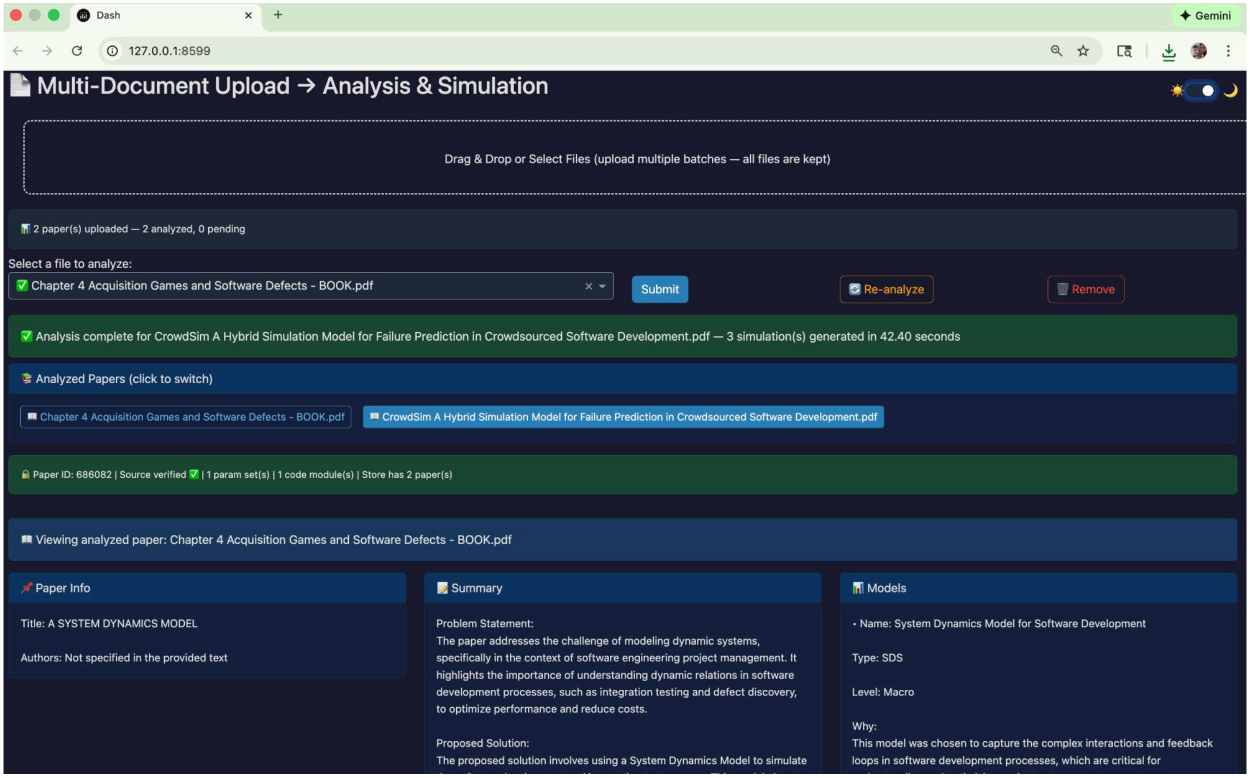

Figure 18. The Simulator Generator's Multi-Document Upload Interface

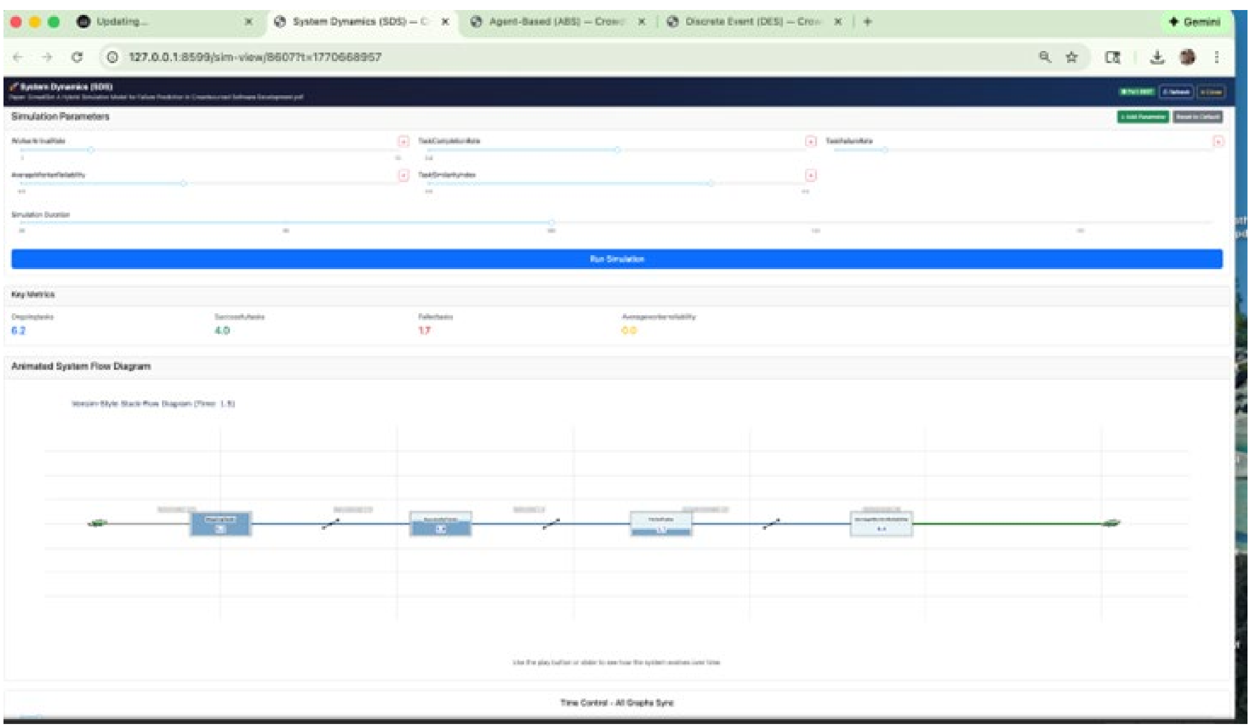

Figure 19. A Generated System Dynamics Model Simulation

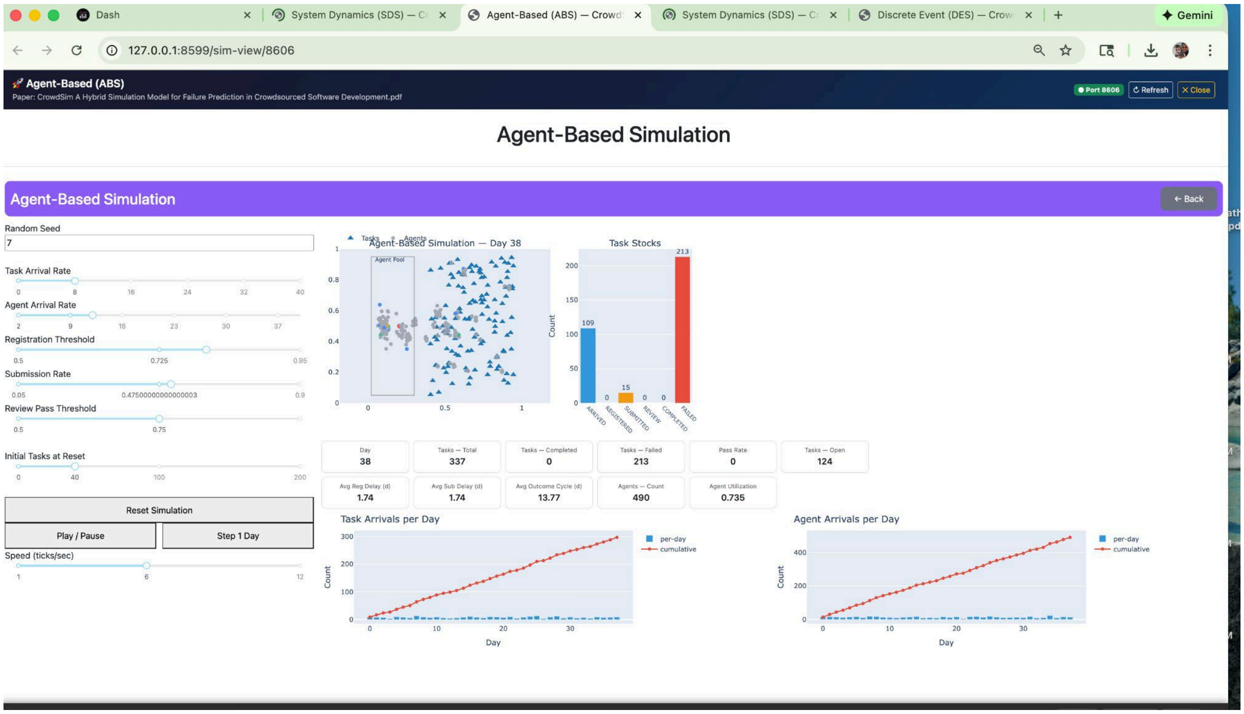

Figure 20. Agent-Based Model Simulation

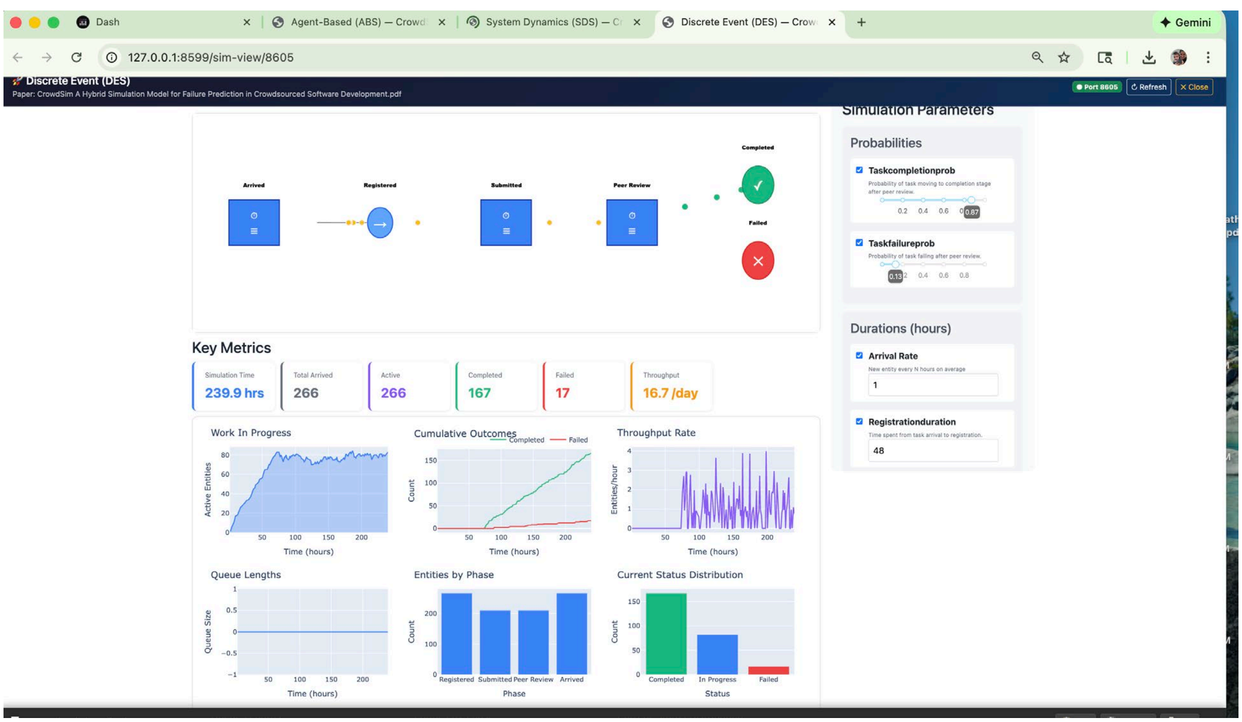

Figure 21. Discrete Event Simulation

The PTL's Database

The other key component of the PTL is the database. ***Table 8*** provides the data tables and purpose by domain in our initial release.

Table 8. PTL Database Inventory by Domain

Domain	Table	Primary Purpose
Session Core	sessions	Master session metadata and UI flags
Session Core	Session_history	Ordered user/assistant history
Session Core	Session_params	Key-value session parameters
Session Core	Session_scenario	Scenario payload and research framing
Session Outputs	Session_research	Research paper rows and metadata
Session Outputs	session_viz	Visualiztaion specs and order
Session Outputs	Session_events	Append-only event timeline
Model Lifecycle	Generated_models	Generated model registry
Model Lifecycle	Session_models	Session-model link map
Model Lifecycle	Model_runs	Model execution runs and status
Model Lifecycle	Model_assets	Model-generated artifacts

The design strengths for the current schema include session-rooted coordination boundaries, append-only event logging via "session_events," flexible JSON-capable payload fields, and pluggable search backends throughout. The key rationale is controlled extensibility: new agent features can be introduced at the orchestration layer while the persistence contract remains stable enough for migration and audit, as shown in Figure 8.

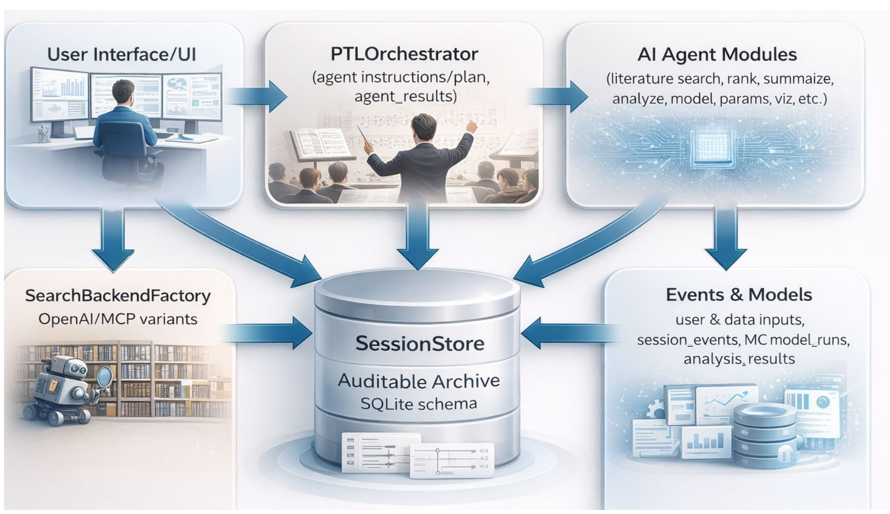

Figure 22. Multi-Agent Persistence and Auditable Coordinated Information Flow

The Enterprise Requirements & Acquisition Model

AIRC obtained from The Aerospace Corporation the latest version of the Enterprise Requirements & Acquisition Model (ERAM) based on Lt Col (Ret.) Robb Wirthlin’s MIT PhD dissertation (Wirthlin et al., 2011). We plan to use ERAM to help jump start our database of models contained within the PTL to build new models that can be compared against. Further, we have already used it to benchmark a prototype of our agent-based approach.

Figure 9 shows a section out of the ERAM 2.3 model, which used the ExtendSIM tool. ERAM contains events for the entire Planning, Programming, Budgeting, and Execution (PPBE) process with constraints and representations of the progression from Pre-A phase through Phase C with funding- and governance-driven delays. Programs entering Pre-A are “gated” by funding availability and Acquisition Category (ACAT) classification, which then determines the magnitude of administrative and review delays. Additional PPBE effects are captured through preparation activities (e.g., acquisition panels, Request for Proposal coordination, etc.) and the possibility of rework loops when decisions or funding alignment fail.

We performed a preliminary comparison of the model’s ExtendSIM outputs to outputs from our AI agents decomposing ERAM 2.3 into a statistically equivalent python model was compared using a Monte Carlo run on a distribution of 500 programs across ACAT I, II, and III program types with success rate as the core validation metric. The approach used LLMs to “mimic” a distribution inside the “black box activity” in an attempt to match the output of the model. This comparison yielded an overall accuracy of 96.6% across all key metrics and demonstrates that our approach is sound, but a lot more work is needed to generalize the ERAM 2.3 model for our full needs.

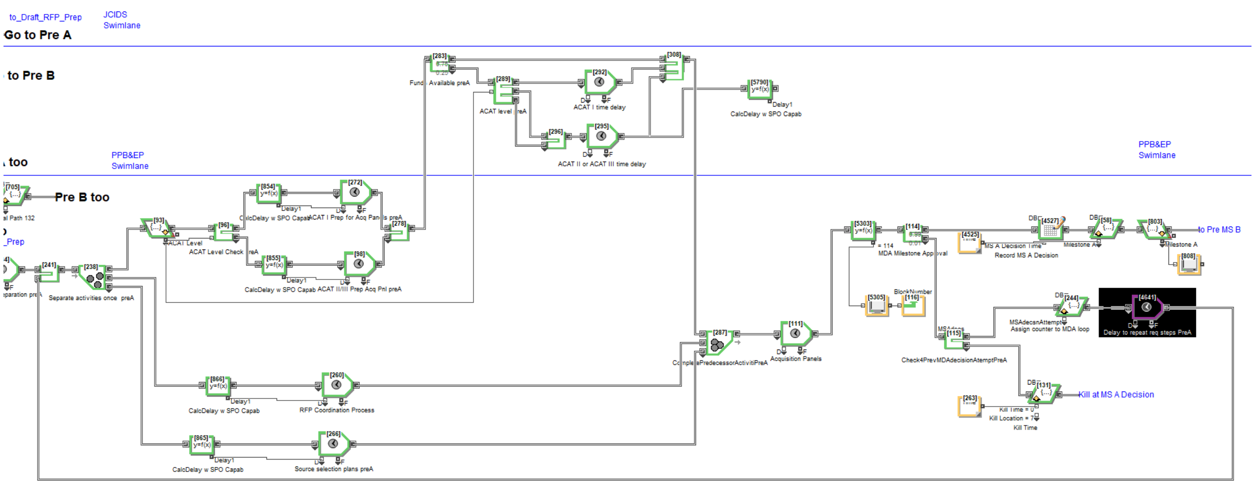

Figure 23. Aerospace’s Lt Col (Ret.) Rob Wirthlin’s MIT PhD Model, ERAM

Table 9 provides the results of the Original ERAM 2.3 model and the LLM agent generated SimPy version.

Table 9. Test Set Performance (Validation)

Metric	Original	SimPy	Delta (%)
Success Rate	53.0%	48.9%	4.1%
ICD (Average Time)	843 days	859 days	2.0%
Cost Growth	1.45x	1.39x	4.0%

Future Plans and Discussion

Decision-makers will soon have a tool that is currently under development, which promises to revolutionize the field of policy analysis using agentic AI. AIRC is well on its way to building a tool for the DoW and its stakeholders to use to try and answer perplexing policy questions. The tool is, however, built in a manner that allows it to adapt to any policy area.

Even though the PTL is still in it is infancy, the ability to research academic literature and create the standard models used by policy researchers is a first step towards a Multi-Level

Computational Policy Simulation that grows over time to provide decision-makers with a capability for doing "*What if?*" policy impact studies looking for unintended consequences.

AIRC was able to identify an Agentic AI startup company (Anadyr Horizon Inc.) with an existing modeling capability that could adapt their analytical predictive intelligence tools, and recently, our OSW leadership facilitated a meeting with CORAS.ai, which has a complimentary effort underway to build a policy document analysis tool for the Navy (see Appendix C).[4]

We anticipate that our efforts will be coordinated in the near future.

Appendix A. Rouse's Ten-Step Methodology

We provide as is, a direct quote of Rouse's ten-step methodology (Rouse 2014, pp. 36-38) for easy reference.

Step 1: Decide on the Central Questions of Interest

The history of modeling and simulation is littered with failures of attempts to develop models without clear intentions in mind. Models provide means to answer questions. Efforts to model socio-technical systems are often motivated by decision makers' questions about the feasibility and efficacy of decisions on policy, strategy, operations, etc. The first step is to discuss the questions of interest with the decision maker(s), define what they need to know to feel that the questions are answered, and agree on key variables of interest.

Step 2: Define Key Phenomena Underlying These Questions

The next step involves defining the key phenomena that underlie the variables associated with the questions of interest. Phenomena can range from physical, behavioral, or organizational, to economic, social or political. Broad classes of phenomena across these domains include continuous and discrete flows, manual and automatic control, resource allocation, and individual and collective choice. Mature domains often have developed standard descriptions of relevant phenomena.

Step 3: Develop One or More Visualizations of Relationships Among Phenomena

Phenomena can often be described in terms of inputs, processes, and outputs. Often the inputs of one phenomenon are the outputs of other phenomena. Common variables among phenomena provide a basis for visualization of the set of key phenomena. Common visualizations methods include block diagrams, IDEF, influence diagrams, and systemigrams.

Step 4: Determine Key Tradeoffs That Appear to Warrant Deeper Exploration

The visualizations resulting from Step 3 often provide the basis for in-depth discussions and debates among members of the modeling team as well as the sponsors of the effort, which hopefully includes the decision makers who intend to use the results of the modeling effort to inform their decisions. Lines of reasoning, perhaps only qualitative, are often verbalized that provides the means for immediate resolution of some issues, as well as dismissal of some issues that no longer seem to matter. New issues may, of course, also arise.

Step 5: Identify Alternative Representations of These Phenomena

Computational representations are needed for those phenomena that will be explored in more depth. These representations include equations, curves, surfaces, process models, agent models, etc. – in general, instantiations of standard representations. Boundary conditions can affect choices of representations. This requires deciding on fixed and variable boundary conditions such as GDP growth, inflation, carbon emissions, etc. Fixed conditions can be embedded in representations while variable conditions require controls such as slider bars to accommodate variations – see Step 9.

Step 6: Assess the Ability to Connect Alternative Representations

[4] For more information about Anadyr Horizon, we refer you to CORAS.ai's website.

Representations of phenomena associated with tradeoffs to be addressed in more depth usually require inputs from other representations and produce outputs required by other representations. Representations may differ in terms of dichotomies such as linear vs. nonlinear, static vs. dynamic, deterministic vs. stochastic, continuous vs. discrete, and so on. They may also differ in terms of basic assumptions, e.g., Markov vs. Non-Markovian processes. This step involves determining what can be meaningfully connected together.

Step 7: Determine a Consistent Set of Assumptions

The set of assumptions associated with the representations that are to be computationally connected need to be consistent for the results of these computations to be meaningful. At the very least, this involves synchronizing time across representations, standardizing variable definitions and units of measures, and agreeing on a common coordinate system or appropriate transformations among differing coordinate systems. It also involves dealing consistently with continuity, conservation, and independence assumptions.

Step 8: Identify Data Sets to Support Parameterization

The set of representations chosen and refined in Steps 5-7 will have parameters such as transition probabilities, time constants, and decay rates that have to be estimated using data from the domain(s) in which the questions of interest are to be addressed. Data sources need to be identified and conditions under which these data were collected determined. Estimation methods need to be chosen, and in some cases developed, to provide unbiased estimates of model parameters.

Step 9: Program and Verify Computational Instantiations

To the extent possible, this step is best accomplished with commercially available software tools. The prototyping and debugging capabilities of such tools are often well worth the price. A variant of this proposal is to use commercial tools to prototype and refine the overall model. Once the design of the model is fixed, one can then develop custom software for production runs.

The versions in the commercial tools can then be used to verify the custom code. This step also involves instantiating interactive visualizations with graphs, charts, sliders, radio buttons, etc.

Step 10: Validate Model Predictions, at Least Against Baseline Data

The last step involves validating the resulting model. This can be difficult when the model has been designed to explore policies, strategies, etc. for which there inherently is no empirical data. A weak form of validation is possible by using the model to predict current performance with the “as is” policies, strategies, etc. In general, models used to explore “what if” possibilities are best employed to gain insights that can be used to frame propositions for subsequent empirical study.

Summary

The logic of the ten-step methodology can be summarized as follows, with emphasis on Steps 1-7:

- Define the question(s) of interest
- Identify relevant phenomena
- Visually compose phenomena
- Identify useful representations
- Computationally compose representations

Note that this logic places great emphasis on problem framing and formulation. Deep computation is preserved for visually identified critical tradeoffs rather than the whole problem formulation. Steps 8-10 of the methodology are common to many methodologies.

Not all problems require full use of this ten-step methodology. Often visual portrayals of phenomena and relationships are sufficient to provide the insights of interest. As just noted, such views are also valuable for determining which aspects of the problem should be explored more deeply.

Appendix B. Additional selected 4+1 Architectural views for the MCPS

In this appendix we provide the remainder of the primary architectural views of the Multi-Level Computational Model.

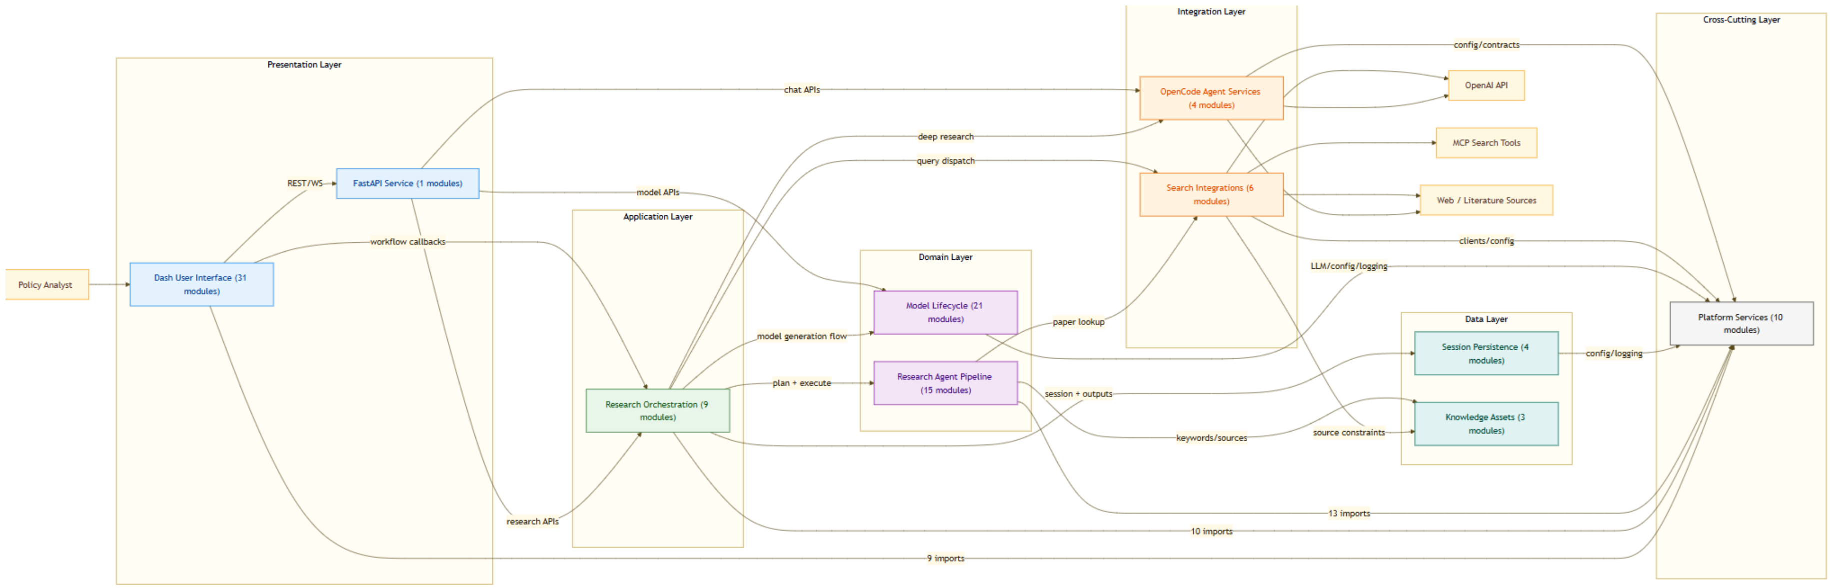

Figure 24. Logical View: Component Diagram

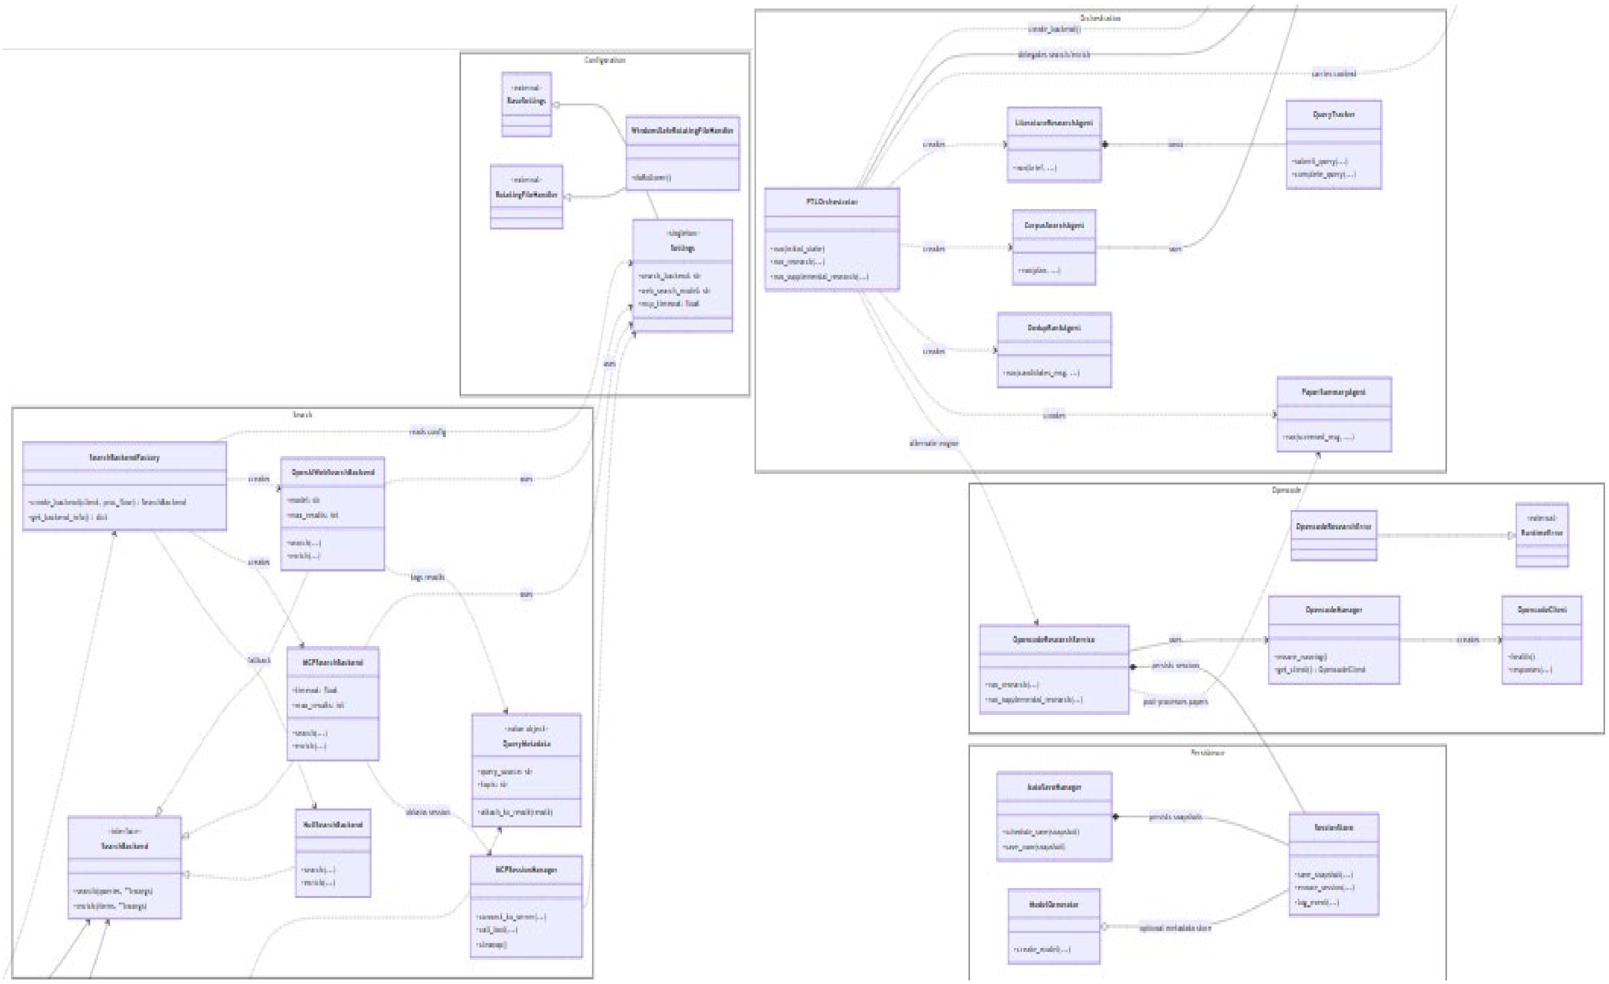

Figure 25. Logical View: Class Diagram – High Level

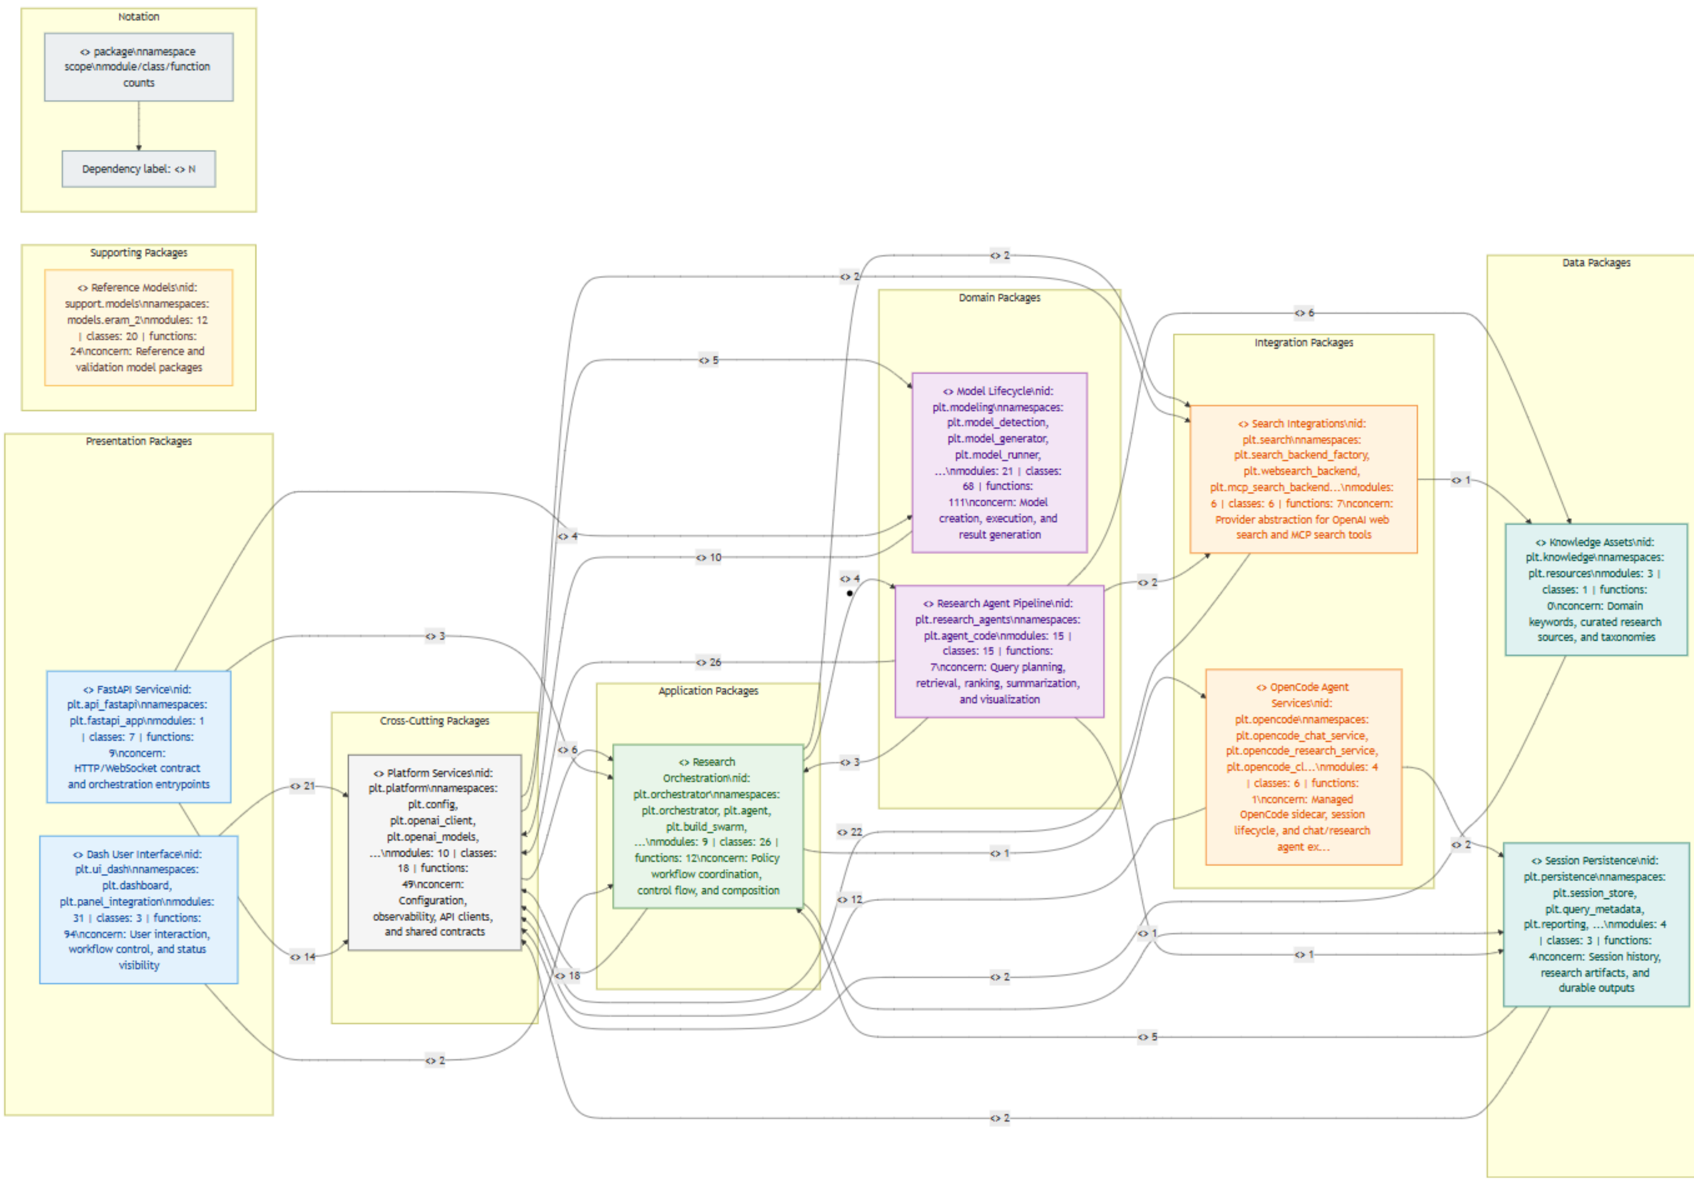

Figure 26. Development View: PTL Package Diagram

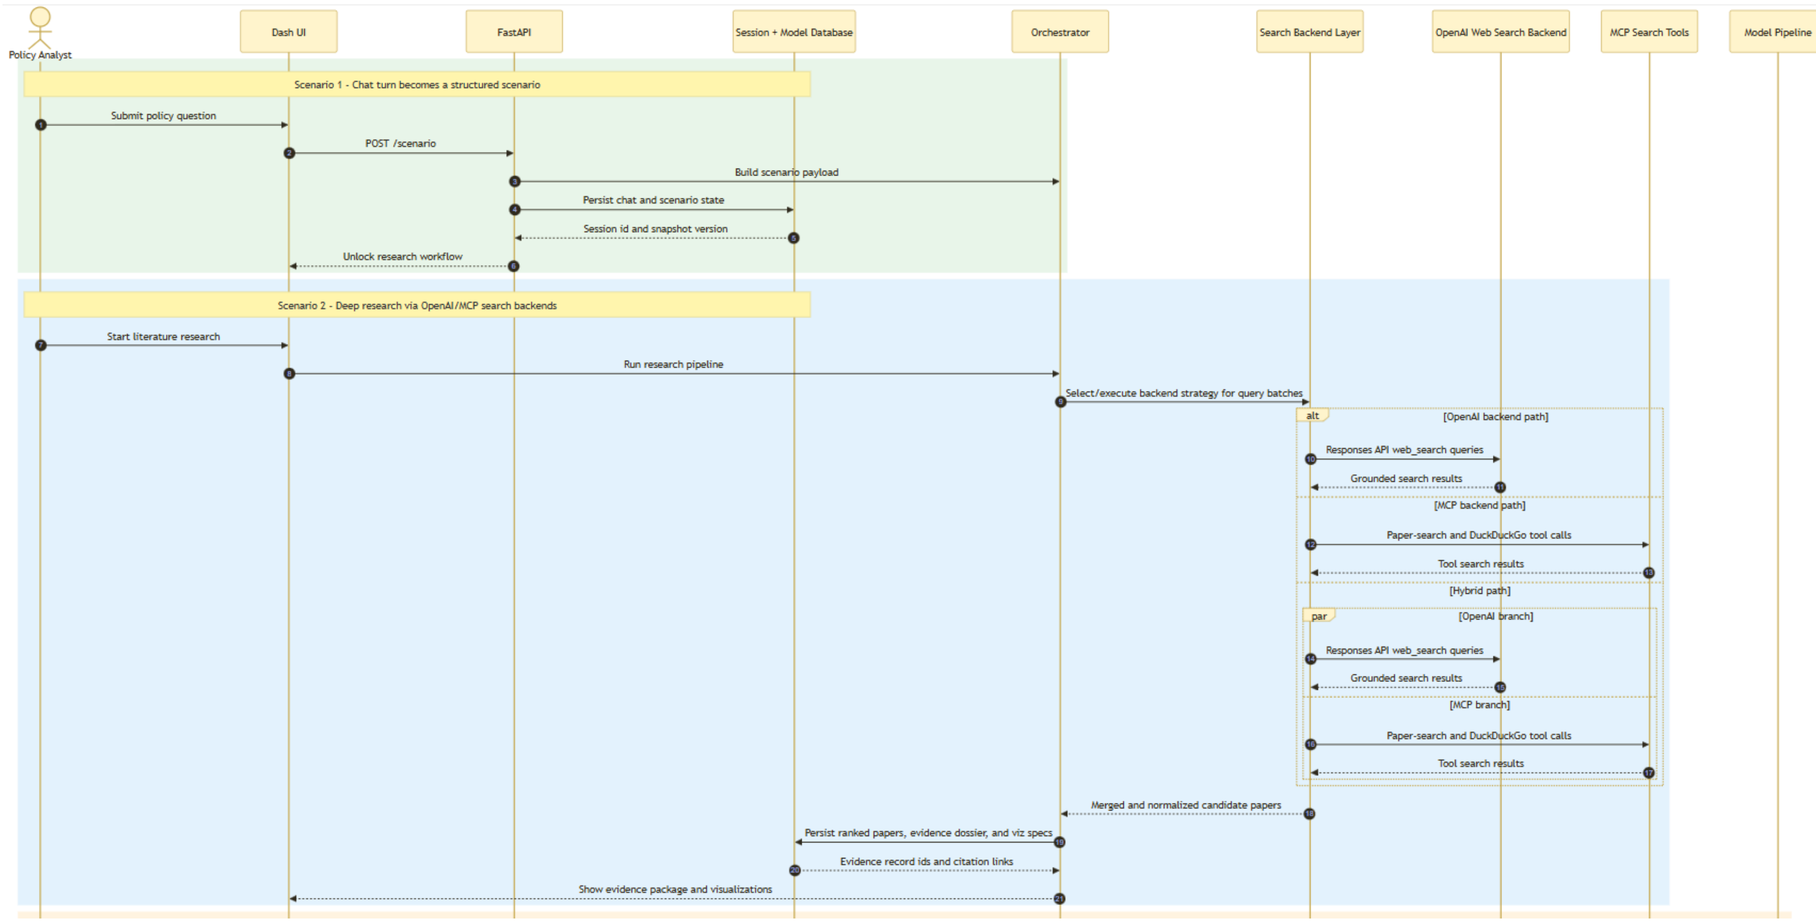

Figure 27. Scenarios: Standard Agent Pipeline Version (Top Half)

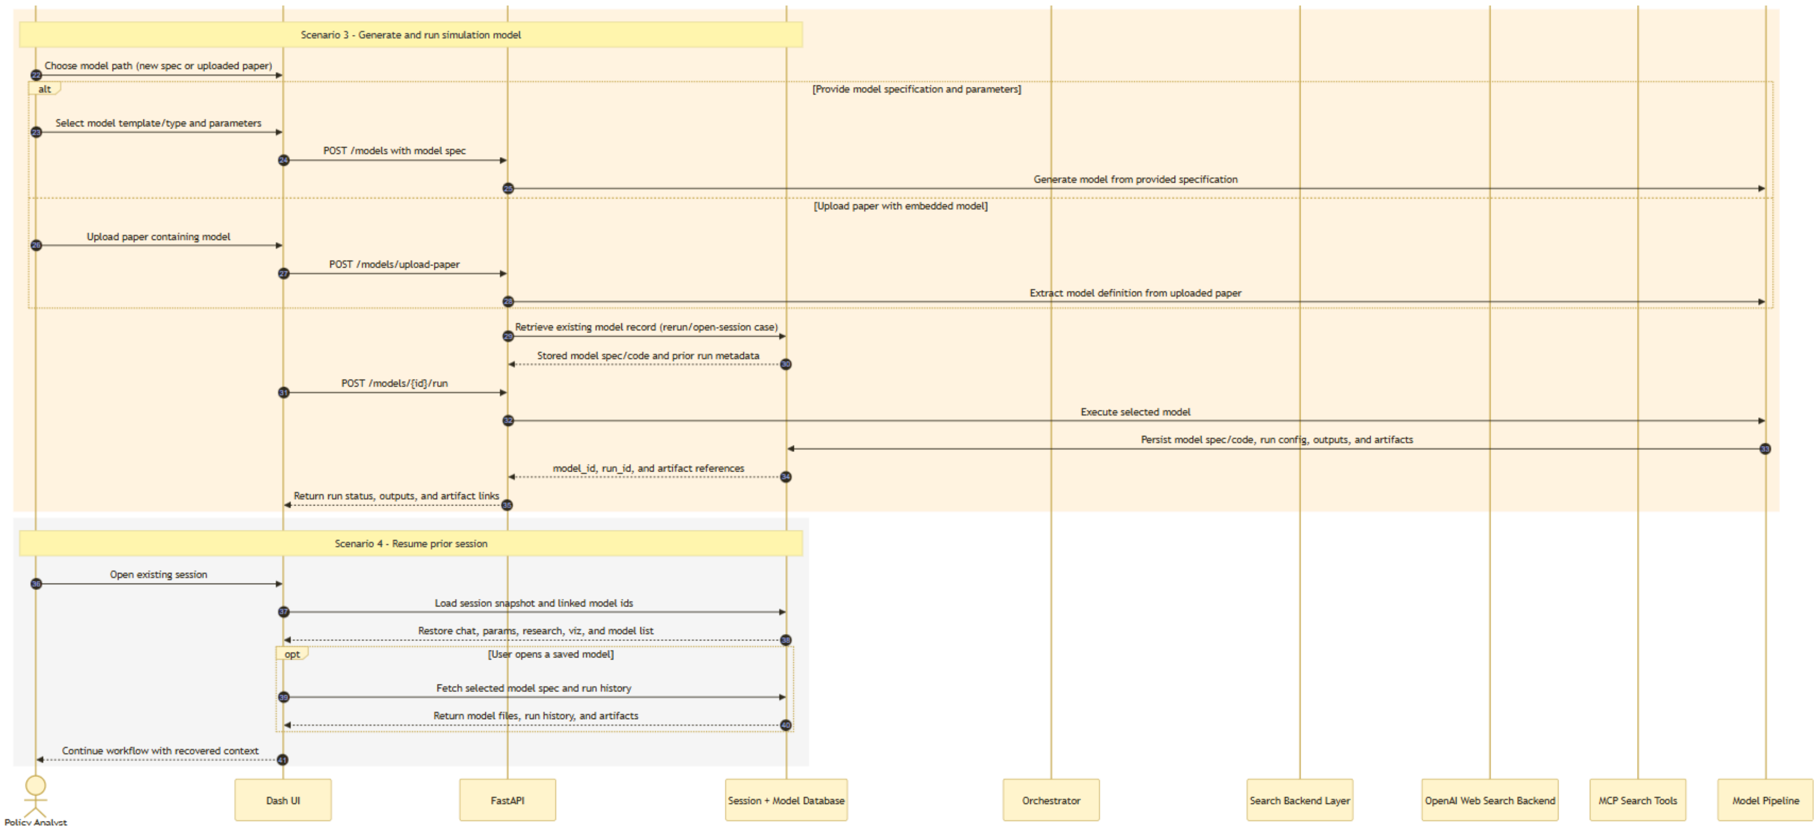

Figure 28. Scenarios: Standard Agent Pipeline Version (Bottom Half)

Appendix C. Agentic Systems Intelligence (ASI)

Simulation Infrastructure for Modeling Decision Dynamics in Geopolitical Crises

Provided courtesy of Anadyr Horizon Inc.

Problem Context

Strategic crises increasingly unfold across interconnected political, military, economic, and cyber systems. Decision-makers must navigate these environments under time pressure, incomplete information, and conflicting institutional incentives. Traditional analytic approaches—including expert judgment, static forecasting models, and episodic tabletop exercises—often struggle to capture how decisions propagate across complex systems and generate cascading effects.

Recent research in crisis simulation and negotiation systems highlights the importance of modeling **dynamic, multi-actor decision environments** rather than isolated events.

Technical Approach

Anadyr Horizon's proprietary technology, **Agentic Systems Intelligence (ASI)**, models geopolitical crises as networks of interacting decision-making agents embedded within strategic systems.

Each agent represents a real-world actor whose behavior is parameterized using historical data, behavioral research, and multi-level constraints. The system integrates agent-based modeling, probabilistic simulation, and scenario stress testing to generate large numbers of potential decision trajectories under defined crisis conditions.

Outputs and Applications

ASI enables analysts to explore how political signaling, military actions, sanctions, or misperceptions may propagate through a system of actors. **North Star**, Anadyr's ASI-powered software platform, generates probability-weighted escalation pathways, identifies potential tipping points and intervention windows, and produces scenario trees illustrating alternative crisis trajectories.

Applications include crisis planning, escalation analysis, red-team exercises, and geopolitical risk modeling. ASI does not attempt deterministic prediction. Instead, it produces probabilistic decision trajectories designed to inform expert judgment, strategic planning, and policy analysis.

Acronyms and Abbreviations

ABM	Agent-Based Model
AFRL	Air Force Research Laboratory
AI	Artificial Intelligence
AIRC	Acquisition Innovation Research Center
ASI	Agentic Systems Intelligence
C2	Command and Control
DES	Discrete Event Simulation
DLA	Defense Logistics Agency
DoD	Department of Defense
DoW	Department of War

DPCAP	Defense Pricing, Contracting, and Acquisition Policy
DPI	Dots Per Inch
DTIC	Defense Technical Information Center
ERAM	Enterprise Requirements & Acquisition Model
JSON	JavaScript Object Notation
LLM	Large Language Model
MCPS	Multi-Level Computational Policy Simulator
MIT	Massachusetts Institute of Technology
MITRE	MITRE Corporation
OV-1	Operational View
PTL	Policy Test Lab
RAND	RAND Corporation
R&D	Research and Development
RDT&E	Research, Development, Test, and Evaluation
SERC	Systems Engineering Research Center
SDM	System Dynamics Model
UARC	University-Affiliated Research Center
UI	User Interface
USD	Under Secretary of Defense
USD(A&S)	Under Secretary of Defense for Acquisition and Sustainment

Acknowledgements

We would like to thank SERC's Mr. Tom McDermott and Ms. Valerie Sitterle for guidance, as well as Ms. Stephanie Halcrow, Dr. Arun Seraphin, and Mr. Moshe Schwartz for the numerous conversations that led the team to investigate the use of the rapidly evolving Agentic AI (Large Language Model–based agents) to solve this challenging problem. We also gratefully acknowledge Dr. Mo Mansouri and Rai Singh Bindra for obtaining results from ERAM, The Aerospace Corporation's Andy Schickling for his support in obtaining ERAM, and Dr. Dan Houston for his help with running it.

Importantly, we want to thank our OSW(A&S)/API sponsors Mr. David Cadman, Ms. Anne Twist, and Ms. Nickee Abbott for their support.

References

DoW. (2025). *Acquisition transformation strategy: Rebuilding the arsenal of freedom*. https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/ACQUISITION-TRANSFORMATION-STRATEGY.PDF

Kannan, H., Rouse, W., Merchant, N., Salado, A., Son, Y., & Szajnfarber, Z. (2022). *The Acquisition Innovation Research Center: Innovation for digital transformation and policy analytics*. Acquisition Innovation Research Center. https://acqirc.org/publications/research/the-acquisition-innovation-research-center-innovation-for-digital-transformation-and-policy-analytics/

Rouse, W. B. (2014). *Visualization of phenomena: Policy flight simulators, modeling and visualization of complex systems and enterprises* (1st ed.). Wiley.

Rouse, W. B., & Bodner, D. A. (2012). *Multi-level modeling of socio-technical systems* (Technical Report No. SERC-2012-TR-020-1). Systems Engineering Research Center. https://sercproddata.s3.us-east-2.amazonaws.com/technical_reports/reports/1524587680.pdf

Rouse, W. B., & Bodner, D. A. (2013a). *Multi-level modeling of socio-technical systems: Phase 1* (Technical Report No. SERC-2013-TR-020-2). Systems Engineering Research Center. https://sercproddata.s3.us-east-2.amazonaws.com/technical_reports/reports/1524587866.pdf

Rouse, W. B., & Bodner, D. A. (2013b). *Multi-level modeling of socio-technical systems: Volume 2* (Technical Report No. SERC-2013-TR-020-3). Systems Engineering Research Center. https://sercproddata.s3.us-east-2.amazonaws.com/technical_reports/reports/1530802703-SERC-2013-TR%20020-3%20Multi%20Level%20Modeling%20Socio%20Technical%20SystemsVolume%202_RT%20044.pdf

Rouse, W. B. Naylor, M. D. PhD, Yu, Z., Pennock, M. J., Hirschman, K. B., Pauly, M.V., & Pepe, K. M. (2019, July–August). Policy flight simulators: Accelerating decisions to adopt evidence-based health interventions. *Journal of Healthcare Management*, *64*(4): 231–241. https://www.doi.org/1097/JHM-D-18-00114

Wirthlin, Joseph R., Houston, Daniel X., & Madachy, Raymond J. (2011). Defense acquisition system simulation studies. *Proceedings of the 2011 International Conference on Software and Systems Process*. https://hdl.handle.net/10945/60714

THIS PAGE INTENTIONALLY LEFT BLANK

PANEL 18. FISCAL AGILITY AND DIGITAL PRECISION: SCALING MODERNIZATION THROUGH NEW FUNDING AND ENGINEERING MODELS

Thursday, May 7, 2026, 1315 – 1430 ET (1015 - 1130 PT)

Panel Summary:

Accelerating warfighting capabilities requires a departure from rigid, hardware-centric fiscal and engineering processes that create a "valley of death" between modernization and operational sustainment. This panel presents two high-impact solutions to this challenge: the reinterpretation of the Defense Working Capital Fund (DWCF) as a capability-based revolving fund to ensure self-sustaining innovation, and the implementation of Digital Technical Data Packages (TDPs) to drastically reduce the cost and time of combat system integration. By linking operational demand with market-aligned fiscal loops and replacing costly live-fire tests with virtual validation, these research initiatives provide a framework for a more responsive and accountable acquisition ecosystem.

Chair: Terry Blake, VADM USN (ret.), Professor of the Practice, Naval Postgraduate School; former Deputy Chief, Naval Operations, Integration of Capabilities and Resources

Discussant: Tom Steffens, Professor of the Practice, Department of Acquisition, Finance and Manpower, Naval Postgraduate School

Panel Presenters:

> **Leveraging the Working Capital Fund Model to Scale and Sustain Innovation –** *Major Rachel Kim, Joint Program Executive Office Armaments and Ammunition*
>
> **Quantifying Return on Investment in Digital Technical Data Packages for Combat Systems: A Lifecycle Perspective –** *John Fiore, Vice President of Strategic Growth, Kitty Hawk Technologies*

Terry Blake, VADM USN (ret.)— served more than 37 years in the United States Navy before retiring in February 2013. He most recently served as the Deputy Chief of Naval Operations for Integration of Capabilities and Resources (OPNAV N8).

As the Navy's Chief Financial Officer, he was charged with planning, programming and executing the Navy's Budget. He also served in numerous positions in the Pentagon including Deputy Assistant Secretary of the Navy (Financial Management & Comptroller) and Deputy Director Resources and Acquisition (Joint Staff J8). His sea commands included a Destroyer, an AEGIS Cruiser, and a Carrier Strike Group.

Since retiring from the Navy, VADM Blake has served as an independent consultant on a number of projects both in the defense and commercial sectors.

He graduated from the U.S. Naval Academy with a Bachelor of Science degree in Political Science, a Master of Science degree in Finance from the Naval Postgraduate School and a Master of Science degree in National Security from the National War College. Additionally, he completed the Seminar XXI program in International Relations from MIT.

Tom Steffens—has served within the Department of Defense for over 40 years to include over 27 Years as an Army Financial Management Officer. He was appointed to the Senior Executive Service in May 2013 and completed his Civilian service as Deputy Chief Financial Officer, Department of Defense from 2022-2024. In this position, Tom was the senior advisor to the Under Secretary of Defense (Comptroller)/Chief Financial Officer and other key Defense leaders on all issues concerning compliance with Congressional mandates and development of DoD initiatives designed to improve Departmental Financial Operations and integrity of Financial Reporting. Previously he served as the Chief Financial Officer for the US Army Corps of Engineers (USACE) for over six years supporting the budgeting, financial operations and reporting of nearly $40 billion in annual outlays supporting our Nation's critical Civil Works & Disaster Response and Military Construction Programs executed by over 36,000 personnel, obtaining unmodified (clean) annual financial statement audit opinions during his entire tenure. Prior to that he was a Director of Accountability and Audit Readiness, Office of the Assistant Secretary of the Army (Financial Management and Comptroller) where he led preparations for the first Army-wide audit of Financial Statements by an Independent Public Accountant.

Tom retired from the Army at the rank of Colonel serving at every leadership level through Battalion Commander with additional assignments ranging from Assistant Professor of Economics, National Defense University; Chief of PPBE Integration, Assistant Secretary of the Army, Financial Management & Comptroller; CFO, U.S. European Command; and Finance Detachment Commander in support of 1st Special Forces Operational Detachment-Delta.

Tom was a recipient of the Presidential Rank Award in 2022 and has previously received the DoD Distinguished Civilian Service Medal and both the Distinguished Civilian Service Award and Superior Civilian Service Award while with Department of the Army. Among his military awards are the Bronze Star Medal while deployed as Battalion Commander in Iraq, the Defense Superior Service Medal and the Legion of Merit (2nd Award). He has both a Certified Government Financial Manager and Certified Defense Financial Manager.

A native of New Jersey, Tom earned his bachelor's degree in accounting from Saint Peter's University, Jersey City in 1985. He received his MBA as a graduate of the Defense Comptrollership Program (then The Army Comptrollership Program) at Syracuse University in 1993 and M.S. In National Security Strategic Studies from the Naval War College in 2007.

Leveraging the Working Capital Fund Model to Scale and Sustain Innovation

Major Rachel Kim, USA—Army Acquisitions Program Manager, Joint Program Executive Office Armaments and Ammunition.

Abstract

The Department of War is undergoing an ambitious acquisition transformation, but the underlying fiscal architecture that governs how capability is funded, scaled, and sustained has remained unchanged since 1949. This paper argues that the Working Capital Fund (WCF), authorized under 10 U.S.C. § 2208, is an underutilized instrument of incentive architecture that, if reinterpreted, can deliver the patient capital that the defense innovation ecosystem requires. Drawing on structural comparisons between the current Defense Working Capital Fund (DWCF) and commercial working capital fund models, the paper proposes an Innovation Working Capital Fund (IWCF) that retains every element required by the break-even mandate and the Financial Management Regulation while replacing the input-based denominator (direct labor hours) with an output-based denominator (availability units), recognizing government-contributed assets through four explicit equity offsets, embedding an Innovation Premium within the rate under existing § 2208(c) and § 2208(e)(1) authority and assigning the Portfolio Acquisition Executive the role of fund manager.

Keywords: working capital fund, defense acquisition, sustainment, innovation capital, Portfolio Acquisition Executive, gainsharing, Planning, Programming, Budgeting, and Execution (PPBE) reform

Background

The defense acquisition enterprise is undergoing an ambitious structural realignment since the Goldwater-Nichols Act of 1986. The 2025 Acquisition Transformation Strategy, which redesignates the Defense Acquisition System as the Warfighting Acquisition System, commits the Department to aggressively prioritize the timely and urgent delivery of operations capabilities to the warfighter through faster decision cycles, commercial-first sourcing, and warfighter-driven requirements (DoD, 2025a; Executive Order No. 14265, 2025; Secretary of War, 2025). The Army's May 2025 reorganization established six Portfolio Acquisition Executives (PAE), consolidating requirements, contracting, testing, sustainment, and international sales under decentralized decision authorities to manage programs across traditional boundaries (DefenseScoop, 2025; Roque, 2025). Major defense innovation reform has crowded in more capital, pathways, and competition for rapid prototyping and fielding. The Office of Strategic Capital's (OSC's) credit program, which provides loans and loan guarantees across 31 Covered Technology Categories ranging from $10 million to $150 million and total loan authority of $984 million through FY2026, was designed to attract and scale private investment in potentially overlooked market segments that support critical technology development (National Defense Authorization Act [NDAA], 2023). The Department has concurrently embraced Middle-Tier Acquisition and Other Transaction Authority (OTA) and propelled speed by streamlining Commercial Solutions Opening (CSO) and Commercial Off-the-Shelf (COTS) acquisition through rapid vehicles such as Undefinitized Contract Actions (UCA) and Defense Other Transaction Consortium (DOTC) structures.

While these efforts have bridged various parts of the "valley of death," the underlying funding architecture remains unchanged. The Commission on Planning, Programming, Budgeting, and Execution (PPBE) Reform concluded that the PPBE process is structurally incompatible with the speed, flexibility, and portfolio-level thinking that modern acquisition requires (Congressional Research Service [CRS], 2024a; PPBE Reform Commission, 2024).

The Commission issued 28 recommendations across five critical areas, including biennial budgeting for appropriate accounts, restructuring the budget around major capability activity areas rather than lifecycle phases, consolidating budget line items, addressing the challenges with colors of money, and authorizing the Department to start or accelerate programs using budget authority under continuing resolutions if approved by defense committees. It identified a single systemic problem: The PPBE architecture cannot respond to new information faster than the Program Objective Memorandum cycle allows, cannot redeploy capital from underperforming programs to overperforming opportunities within the fiscal year, and cannot make the multi-year financial commitments that both the industrial base and private capital markets require to make rational investment decisions.

The FY2026 NDAA addressed some elements, such as the PAE authority, but the fundamental fiscal architecture remains intact: annual appropriations, classified by color and purpose, expiring at the fiscal year boundary, allocated through a two-year budget cycle that produces funding decisions 18 to 24 months before execution and cannot be revised without triggering a laborious reprogramming process (Purpose Statute, n.d.; Anti-Deficiency Act, n.d.; Reprogramming of Funds, n.d.). Congress has not passed a defense appropriations bill on time since 1997, meaning that for nearly three decades, every fiscal year has begun under a continuing resolution that restricts new starts, prevents production rate increases, and freezes program obligations at prior-year levels (Duffey, 2026; GAO; SC World, 2026). Budgeting reform has not kept up with acquisition reform in the same manner, exposing the gap between reform language and budget reality (GAO, 2024).

The empirical evidence that the current reforms have not addressed the core problem is documented across four years of successive GAO reports. The June 2025 Weapon Systems Annual Assessment found that cost growth continues to accelerate despite reform: Combined total estimates for 30 major defense acquisition programs grew by $49.3 billion in a single year. Programs using the MTA pathway entered with low technology maturity and experienced the same schedule delays as the traditional path they were meant to replace. Of seven former MTA programs the GAO (2025a) reviewed, none were ready for production or fielding when the MTA effort ended.

This “ditch of death” exists at the transition from production to sustainment (Federal News Network, 2025; MITRE, 2025). The GAO reported that Operations and Maintenance (O&M) dollars account for approximately 70% of a weapon system’s total lifecycle cost. Of 16 weapon systems the GAO assessed in FY2022, nearly half had sustainment cost growth above the thresholds identified in statute. The Department projects to spend approximately $2.4 trillion on its costliest major defense acquisition programs across their operational lifetimes. Despite the enormous capital flows to the supply side, the sustainment enterprise that consumes 70% of lifetime program cost operates with no self-financing innovation mechanism, no performance-based incentive structure, and no instrument to attract private capital at the scale the industrial base modernization challenge requires (DoD, 2022; GAO, 2025b; National Defense Industrial Association [NDIA], 2023). Until sustainment is treated as a coequal pillar of modernization, rather than an afterthought, Army modernization will remain trapped between intent and execution.

Urgent procurement without sustainment financing creates operational dependency on commercial supply chains that neither the government nor its vendors have planned to maintain (GAO, 2024). Once a production contract is awarded, the feedback loop between the warfighter and the industrial base is effectively disincentivized: the contractor, having achieved program of record status, moves from a competitive market into a structural monopoly; the government has made its platform choice; the technical data package (TDP) can be controlled by the original equipment manufacturer; the sustainment relationship is governed by sole-source contracts that

GAO has found account for the majority of post-production spending. If PAEs cannot manage their portfolios like fund managers—with retained earnings, multi-year commitment authority, and the ability to redeploy capital across program boundaries based on performance—then the PAE authority, however ambitious in concept, cannot produce the outcomes the acquisition transformation strategy promises. If appropriations are classified in colors with expiration dates rather than treated as the fungible capital that private sector portfolio managers command, the same mechanical problems will continue to occur: Investment lags because the capital cannot be deployed when the opportunity appears; program overruns because the correction mechanism cannot engage until the next rate cycle; sustainment cost growth because the post-production incentive structure rewards consumption rather than performance. Notwithstanding the concentration of investment across RDTE Budget Activities 6.1 through 6.4—the $71.8 billion RDT&E portfolio that drives the Department's innovation investment—the underlying incentive structure remains oriented toward a singular objective: attainment of program of record status, after which the contractor's financial incentives run directly against the continued modernization and innovation the warfighter requires (DoD, 2025b).

What the defense innovation ecosystem lacks is patient capital investment that can absorb multi-year development cycles, tolerate procurement delays, and remain solvent through the appropriations environment. Drawing on existing authorized programs and currently available statutory authority, this paper argues that the Working Capital Fund (WCF), authorized under 10 U.S.C. § 2208—a $63.8 billion annual instrument that already exists and possesses the structural properties (CRS, 2024b; DoD, 2025c)—would circumvent reliance on volatile appropriation cycles, fund next-generation capability qualification from sustainment cost-recovery revenue, restore competitive discipline through mandatory government ownership of TDP, and create the incentive architecture necessary to crowd in private capital on the supply side. In the interim of larger structural reforms to the PPBE process, this paper reexamines the WCF not as a legacy accounting mechanism but as an underutilized incentive architecture.

How the Current Defense Working Capital Fund Works

The DWCF is the Department's revolving fund mechanism for financing common services and industrial-type activities. Authorized originally in the National Security Act Amendments of 1949 and now codified at 10 U.S.C. § 2208, the fund was established to allow the Department to aggregate demand from multiple military departments and agencies, recover costs through a stabilized billing rate, and smooth cash flow across fiscal years (CRS, 2024b; National Security Act Amendments of 1949, 1949).

It encompasses five activity groups operated by three defense agencies: the Defense Logistics Agency (supply chain management, energy, and document services), the Defense Information Systems Agency (information services), and the Defense Finance and Accounting Service (financial operations). Together, these five activity groups disbursed $63.8 billion in FY2026 and received new orders totaling $62.2 billion (DoD, 2025c).

Figure 1 illustrates a general example of how a DWCF operates and follows the rate-setting methodology described below.

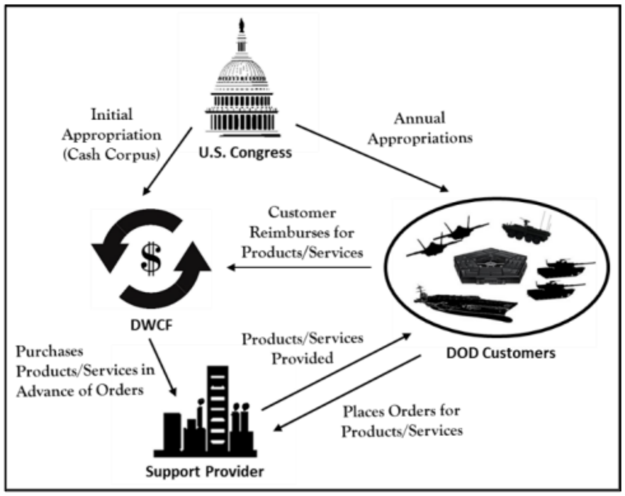

Source: CRS Graphics.

Notes: The process illustrated above is a general example of how a DWCF operates. Variations can exist (e.g., private party customers).

Figure 1. How a DWCF Operates
(CRS, 2024b)

- Ordering activities, such as military departments, defense agencies, and in limited cases allied governments through Foreign Military Sales, submit funded orders to the fund at the published rate, using their O&M or other appropriations.
- The fund executes the work, delivers the service, or ships the supply.
- The fund bills the ordering activity at the published rate and collects the reimbursement.
- At year-end, the difference between actual cost and actual revenue, or the Net Operating Result, accumulates into the AOR balance and is scheduled for recovery or return in a future rate cycle.
- The WCF must submit its Justification Books (J-Books) to Congress annually, which provide complete visibility into each WCF activity's revenues, costs, rate changes, and capital investment program, without requiring separate appropriation for each transaction or increment. Congress retains oversight through transparency rather than transactional control.

Congress appropriates cash or authorizes in-kind inventory transfer to the fund's opening balance. This initial appropriation, a one-time capitalization, combined with any inventory or physical assets transferred in-kind from prior appropriated fund accounts, constitutes the cash corpus. The corpus gives the fund manager working capital to purchase materials, pay labor, and execute the first cycle of work before customer reimbursements arrive. The Office of the Under Secretary of Defense (Comptroller), or the Assistant Secretary of the relevant military department for the service funds, establishes the annual rate schedule 18 to 24 months before the fiscal year in which it takes effect. The rate is set against a projected workload forecast and a projected cost base, incorporating any prior-year Accumulated Operating Result (AOR) correction from the cycle two years back.

The WCF rate is stabilized and published before the fiscal year, held constant throughout the fiscal year regardless of actual cost experience, and corrected only in subsequent rate cycles through the AOR mechanism. This stabilization provides ordering activities with the budget certainty that direct appropriations cannot, absorbing variance in the fund's own balance sheet rather than exposing ordering activities to mid-year cost shocks. The cost of this feature is the 18-to-24-month correction lag between a cost event and its appearance in the rate.

The standard DWCF rate equation consists of four cost-recovery components in the numerator divided by a projected volume of Direct Labor Hours in the denominator:

$$Rate\ (\$/DLH) = \frac{(Direct\ Labor + Indirect\ Overhead + Capital\ Surcharge \pm AOR)}{Projected\ Direct\ Labor\ Hours}$$

$$Direct\ Labor = Civilian\ workforce\ fully-loaded\ compensation$$

$$Indirect = Allocated\ overhead: G\&A, facilities, IT, management$$

$$Capital\ Surcharge = Depreciation\ recovery\ on\ government-owned\ equipment$$

$$AOR = Accumulated\ Operating\ Result: prior-year\ surplus\ (-)\ /\ deficit\ (+)$$

$$DLH = Direct\ Labor\ Hours: the\ input\ unit\ the\ fund\ bills\ against$$

The rate denominator is the projected DLH volume, or the total number of direct labor hours the fund expects to bill in the coming fiscal year, derived from the FTE count and a productive yield assumption. The rate is a single hourly price that ordering activities pay for every hour of depot labor they consume. It is then stabilized for the full fiscal year, meaning customers pay the same rate regardless of whether actual costs come in higher or lower than projected.

This equation has three structural weaknesses. First, the unit of account, DLH, is an input, which means the government buys time, not performance. Second, every component of the numerator is a cost-recovery element. No component rewards performance above baseline, attracts capital investment, or creates a return window. Third, the denominator is a forecast, rather than a validated demand signal. If actual volume falls below projection, the rate under-recovers. If it exceeds projection, the rate over-recovers. Neither outcome triggers any automatic response within the FY.

The break-even mandate, which is the fund's defining legal constraint, requires that rates recover the full costs of operation over a reasonable period. Surplus in excess of working capital requirements must be returned to ordering activities through rate reductions or carried as AOR. Deficits are recovered through rate increases in subsequent years (Working-Capital Funds, n.d.). The AOR is the WCF's self-correction mechanism: the running total of all prior-year surpluses and deficits, which the fund must drive toward zero over the budget cycle. The Net Operating Result (NOR) for any given fiscal year is revenue minus expenses; a positive NOR means the fund over-recovered relative to cost, and a negative NOR means it under-recovered. NOR values accumulate into the AOR balance and are recovered through rate adjustments in subsequent years until the AOR returns to zero.

The FY2026 Budget reveals how this AOR mechanism works in both directions. The Army WCF's Industrial Operations activity group entered FY2026 with an AOR deficit of −$129 million accumulated over three years of under-recovery, driven primarily by workforce reductions under the Army Transformation Initiative that were not priced into the rate at the time it was set. The FY2026 rate was set above the cost floor to zero out this deficit, which means ordering activities are paying a premium rate in FY2026 to recover losses the fund incurred in FY2024 and FY2025. At the same time, the DWCF entered FY2026 with an accumulated prior-year surplus of $5.5 billion, which the fund resolved by writing off $3.05 billion as non-recoverable and deferring $2.56 billion to future years (DoD, 2025c, 2025d). In other words, the ordering activities who funded that $5.6 billion surplus through their O&M appropriations will not receive a rebate.

The Defense-wide WCF cash balance must always remain positive. A DWCF cash balance that goes negative represents an Anti-Deficiency Act violation under 31 U.S.C. § 1341. The Lower Operating Range (LOR) acts as a hard floor that triggers mandatory corrective action. When the DWWCF's cash position of $3.1 billion is compared against the Lower Operating Range of $2.26 billion, the $840 million cushion represents a narrow margin of approximately four days of operating disbursements, and therefore is one of the primary reasons the current rate-setting mechanism cannot absorb much additional stress without triggering either a rate increase, an AOR acceleration, or a corpus replenishment request.

The WCF's cash management operates within two statutory limits: an LOR below which the fund is considered insolvent and an Upper Operating Range (UOR) above which surplus is presumptively subject to offset. Between those two limits, the fund holds its working capital and the stabilization reserve, which absorbs demand volatility and price fluctuation that the annual rate cannot anticipate (Working-Capital Funds, n.d.; DoD, 2025c). The FY2026 DWWCF cash table illustrates both the scale of the fund's cash cycle and the pressure it has been under in recent years.

Table 1. Defense-Wide Working Capital Fund Cash Balance ($K)

DWWCF Cash Balance ($K)	FY24 Actual	FY25 Est.	FY26 Budget
Cash, Beginning of Period	$5,704,207	$4,133,096	$3,652,262
Total Disbursements	**−$59,194,098**	**−$65,046,825**	**−$63,758,247**
Total Collections	+$57,197,601	+$64,563,738	+$63,194,062
Net Cash Loss from Operations	**−$1,571,111**	**−$480,834**	**−$552,216**
Cash, End of Period	$4,133,096	$3,652,262	$3,100,046
Lower Operating Range	$2,586,107	$2,371,318	$2,259,522
Upper Operating Range	$7,183,747	$7,250,500	$7,238,410
Three-year projected cash decline	**−$2,604,161K total**		Cash position declining toward lower bound

Note. Data from the *Defense-Wide Working Capital Fund FY2026 Budget Estimates* (DoD, 2025c).

The DWWCF's cash position has declined by $2.6 billion from the beginning of FY2024 to the end of FY2026, a drawdown approximately equivalent to 18 days of operating disbursements. This rate-setting mechanism cannot adjust fast enough to match cost experience. The stabilization reserve, which was intended exactly for this kind of volatility absorption, has been depleted faster than the AOR correction mechanism can replenish it.

The Capital Investment Program

The WCF finances not only operating expenses but also capital improvements through a Capital Investment Program (CIP), which permits the fund to acquire capital assets (major equipment, automated process improvements, facility modernization, and IT systems) with expected useful lives beyond one year. The CIP is funded through a Capital Surcharge component included in the rate, which recovers depreciation on the acquired assets over their useful lives. This allows the fund to make multi-year capital investments that would otherwise require separate appropriations, and it ensures that the cost of those investments is borne by the ordering activities who benefit from the resulting operational improvements.

The AWCF Industrial Operations CIP grew sharply in FY2026, from $100.3 million in FY2025 to $197.7 million in FY2026. The CIP authority is one of the few existing mechanisms through which the WCF can make forward-looking investment rather than pure cost recovery (DoD, 2025d).

A final structural feature of the current WCF model is that every DWWCF activity group operates on a mandatory-use basis, which means that ordering activities cannot source the services the fund provides from any other provider. The Army, Navy, and Air Force cannot buy their financial accounting services from anyone but the Defense Finance and Accounting Service. The services cannot replace DLA as the supply chain manager for the 5 million consumable items DLA manages. They cannot source their information services outside the Defense Information Systems Agency for the mission sets it operates. Under 10 U.S.C. § 2466, the Army, Navy, and Air Force cannot source more than 50% of their depot-level maintenance from private contractors, effectively mandating use of the organic industrial base financed through the service working capital funds.

The mandatory-use structure preserves the organic industrial base as a strategic resource independent of commercial market availability, maintains readiness capabilities that would atrophy if exposed to normal commercial competition, and ensures continuity of critical services during periods of commercial market disruption. However, the consequence of the structure is that the WCF rate is not a market price, but an accounting allocation. The provider has no competitor offering the same service at a lower rate, and the ordering activity has no alternative supplier to threaten switching to. Whatever financial pressure exists on the provider's cost performance must come from the internal rate-setting process, not from external market discipline. This is the structural environment in which the feedback loop between efficiency and reward either operates or does not.

Characteristics of the Model That Must Be Retained

- The WCF provides multi-year execution across fiscal years and thus is resilient to budget lapses. When Congress fails to pass a defense appropriations bill, which has been an annual occurrence since 1997, the WCF continues to operate because it is funded through this revolving fund. No other defense financial instrument provides this without specific multi-year appropriation authority.
- The WCF aggregates demand across all military departments and defense agencies into single financial relationships with its provider activities. For instance, DLA Supply Chain Management manages supply chains for the Army, Navy, and Air Force simultaneously, enabling purchasing leverage that no individual service could achieve independently.
- The WCF provides price stability for ordering activities who can plan off their respective direct appropriations. Ordering activities budget against a published rate that is held constant for the full fiscal year, thus acting as an insurance function. They can plan against a stable sustainment rate to manage their budget decisions rather than the alternative of appropriated funds, which can be exposed to mid-year cost volatility.
- The WCF produces full-cost visibility. When a program is funded through direct appropriation, the visible cost to the program is the contract price; the overhead costs of the depot that performs the work, the depreciation on the equipment used, and the management burden of the contracting activity are all absorbed into defense-wide overhead that the program manager never sees. When the same program is funded through the WCF, all of those costs appear in the rate, making the true cost of the service visible.
- External Sales Authority (10 U.S.C. § 2208(j) and 10 U.S.C. § 4543). Section 2208(j) authorizes WCF-funded industrial facilities to manufacture and sell articles to persons outside the Department when the purchaser is fulfilling a defense contract and the

solicitation was open to competition. Section 4543 extends this to Army industrial facilities manufacturing ammunition, munitions, or components thereof. For instance, Foreign Military Sales can generate additional revenue, creating a strategic export dividend that reduces the per-unit cost burden on domestic customers.

How the Private Sector Manages Working Capital

In the private sector, working capital is a fundamental operating metric defined as current assets minus current liabilities. Firms actively manage this balance to ensure liquidity while deploying capital productively. Unlike fixed capital investments—plant, equipment, long-term debt—working capital finances the short-cycle operations that keep the enterprise running: raw material procurement, payroll, accounts receivable, and inventory.

A healthy firm manages working capital through the cash conversion cycle (CCC): the elapsed time between cash outflow for inputs and cash inflow from sales. Firms with shorter CCCs—those that collect receivables quickly, turn inventory efficiently, and stretch payables appropriately—require less working capital to sustain the same revenue volume. Lean manufacturers, software-as-a-service companies, and platform businesses often achieve negative CCCs, meaning customers pay before the firm incurs full production costs (CFA Institute, 2024).

In a for-profit model, working capital funds are traditionally managed by setting a billing rate using the following general structure:

Commercial Managed Services Rate

$$\textit{Commercial Managed Services Rate}$$

$$Rate\ (\frac{\$}{output}) = \frac{Cost\ Floor\ +\ Operating\ Margin\ +\ Reinvestment\ Premium\ +\ Risk\ Reserve}{Contracted\ Output\ Volume}$$

Cost Floor = All direct and indirect operating costs.

Operating Margin = 8–15% retained earnings.

Reinvestment Premium = 2–5% (typically).

Risk Reserve = 1–3% contingency.

Output Volume = unit of output.

The commercial fund's rate is constructed to generate four distinct layers of value: cost recovery, operating margin, reinvestment premium, and risk reserve. Cost recovery captures the direct and indirect costs of delivering the service (the equivalent of the DWCF's full cost build-up). The operating margin of 8% to 15% for industrial managed services, or 15% to 25% for IT and professional services, is retained by the fund as earnings available for reinvestment, distribution to shareholders, or accumulation as retained earnings (Boston Consulting Group, 2024a; KPMG, 2023). The reinvestment premium of 2–5% of revenue is the self-financing mechanism that allows the fund to modernize without requiring external capital injection. The risk reserve of 1% to 3% absorbs volume variance and demand fluctuation, analogous to the DWCF's stabilization reserve but retained by the fund rather than subject to customer return.

How Commercial Working Capital Funds Are Governed

The commercial fund's governance structure reflects its portfolio orientation. The fund manager is evaluated on three metrics simultaneously: return on invested capital, customer performance outcomes, and forward capability investment. The three metrics are integrated through the rate and the fund's capital allocation authority. A fund manager who beats the performance benchmark retains margin, which funds the reinvestment premium, which in turn improves next period's performance against benchmark. The feedback loop is immediate, attributable, and compounding.

The General Partner/Limited Partner Model

In more sophisticated structures, the commercial WCF is organized as a closed-end fund with a General Partner (GP) managing capital supplied by Limited Partners (LPs). The GP makes the portfolio allocation decisions: which assets to acquire, what capital improvements to fund, when to exercise governance rights, when to divest. The LPs supply the capital and receive their returns through distributions based on fund performance (Blackstone Group, 2024; Macquarie Asset Management, 2024). The GP earns a management fee (typically 1%–2% of committed capital annually) covering operational expenses and a performance fee (typically 20% of profits above an 8% hurdle rate), or carried interest, that directly aligns the GP's incentives with LP returns.

In a commercial fund, efficiency gains flow through a three-step cycle: the gain appears as margin improvement in the current period; the margin flows to retained earnings or distribution at the fund manager's discretion; retained earnings fund the next round of process improvement investment, which generates the next efficiency cycle. The cycle is self-reinforcing. A commercial depot that achieves a 10% labor productivity improvement captures the full 10% as margin for at least the life of the contract period, deploys 2%–5% into further reinvestment, and distributes the remainder. The next contract cycle, the depot negotiates from a stronger competitive position because it has demonstrated performance.

The defense WCF's break-even mandate converts this cycle into a liability. Every efficiency gain above the cost floor becomes surplus that 10 U.S.C. § 2208(m) requires be returned to customers. The same 10% productivity improvement that a commercial depot captures as margin flows immediately into next year's rate reduction at a defense depot. The depot has nothing left to reinvest. The next round of process improvement requires competing for appropriations in the normal PPBE cycle against every other defense priority.

Incentive Structure: The Profit Motive as Innovation Engine

The defining feature of private-sector WCF management is the alignment between working capital efficiency and shareholder value. Firms that deploy working capital more productively—generating higher return on invested capital (ROIC) per dollar of liquidity deployed—create equity value. This incentive cascades through the organization: Division managers who generate cash from operations earn autonomy and resources; those who consume working capital without revenue justification must explain the gap.

Crucially, this incentive structure rewards innovation that reduces cost, expands revenue, or compresses cycle time. A product manager who modernizes a sustainment offering to command a recurring software subscription fee creates working capital—cash flow that the firm can redeploy. The same logic drives continuous improvement in manufacturing, logistics, and support: Efficiency gains translate directly into capital available for the next investment cycle.

Side-by-Side Comparison: Commercial vs. Defense WCF

Table 2: Structural Comparison: Commercial WCF vs. Defense WCF

Dimension	Commercial WCF	Defense WCF	Structural Consequence
Rate denominator	Output unit (flight hour, available vehicle-day, completed transaction)	Input unit (Direct Labor Hour consumed)	Commercial provider profits from reliability; defense provider profits from consumption
Operating margin in rate	8–15% retained; self-finances innovation and modernization	0%, break-even mandate prohibits retained margin	Commercial fund generates revenue off its base; defense fund generates $0
Reinvestment premium	2–5% range	Not in rate; innovation competes in PPBE appropriations	Commercial: innovation is a rate component. Defense: only with investment dollars
Surplus / AOR treatment	Retained earnings; redeployed at management discretion	Mandatory customer rebate or write-off	Commercial: efficiency creates reinvestment capital. Defense: efficiency creates rebate obligation
Correction cycle speed	Current period; variance absorbed in fund's margin buffer	18–24 months via AOR mechanism	Commercial fund self-corrects; defense fund carries error for two years
Innovation return window	Defined: first-mover margin captured for 2–5 years until competition	None; efficiency gains returned immediately	Commercial fund rewards first mover; defense fund penalizes first mover with rate reduction
Capital leverage ratio	3:1 to 6:1 leverage against contracted revenue	0:1; no multi-year commitment authority	Commercial: $100M equity supports $400–600M of deployed capital. Defense: $100M is $100M
Annual innovation capital generation	3–8% of revenue self-financed; plus 3–5× private leverage crowded in	0% self-financed; 0× private leverage	The gap is an incentive architecture problem

Note: Analysis based on DoD Financial Management Regulation Volume 11B (DoD, n.d.-c), Working-Capital Funds (n.d.), KPMG (2023), BCG (2024a), and Preqin (2024).

The Innovation Capital Gap Quantified

The structural comparison translates into a specific, quantifiable innovation capital gap. Using an example of the Army Abrams sustainment base of approximately $5.1 billion annually, the commercial model produces approximately $255 million in annual self-financed innovation capital, at a conservative 5% operating margin.

The operating margin retained after cost recovery is the primary source of capital that can be reinvested toward innovation. The mechanics of how retained earnings translate into reinvestment can be expressed as:

Commercial: $Innovation\ Capital = Revenue - Cost\ Floor - Required\ Distributions$

$$= EBITDA - Capital\ Charges - Required\ Payout$$

$$= 3 - 8\%\ of\ revenue\ annually, self - financed$$

Applied to a notional $5.1 billion Abrams sustainment base at 5% margin:

Innovation Capital = $5.1B × 0.05 = $255M/yr, self-financed, retained

DWCF: $Innovation\ Capital = (Rate \times Volume) - Cost\ Floor - AOR\ Return$

By break-even mandate: Rate × Volume ≈ Cost Floor

Therefore: Innovation Capital = 0. The gap on an identical $5.1 billion portfolio is $255 million, and over a 20-year program life, this is $5.1 billion in foregone innovation capital.

The gap compounds when private capital leverage is factored in. Commercial infrastructure funds routinely deploy 3:1 to 6:1 leverage against contracted revenue streams, meaning each dollar of retained margin supports $3 to $6 of deployed capital (Deloitte, 2022; Preqin, 2024). The commercial $255 million of retained innovation capital becomes $1.0–1.5 billion of deployed capacity once leveraged through the private market. The DWCF $0 of retained innovation capital supports $0 of deployed private capital because there is no retained fund balance against which to raise co-investment and no multi-year revenue commitment against which private capital can underwrite a return (CFA Institute, 2024; Preqin, 2024). The cumulative gap, on a single program's 20-year lifecycle, is measured in the billions of dollars of foregone industrial modernization.

One notable caveat is the legal interpretation of Section 2208(c)(1) mandating that WCFs shall be charged with the cost of the procurement and qualification of technology-enhanced maintenance capabilities that improve reliability, maintainability, sustainability, or supportability and have, at a minimum, been demonstrated to be functional in an actual system application or operational environment. It is written for industrial-type and commercial-type activities such as supply management, depot maintenance, transportation, and finance, but it is not designed as an investment fund. The statute requires the fund to break even over the long term, and § 2208(m) directs the Department to ensure fund balances do not exceed working capital requirements (Working-Capital Funds, n.d.).

Gainsharing: The Commercial Instrument for Aligning Provider and Customer

A final dimension of the commercial WCF model is the gainsharing instrument. Commercial firms routinely use gainsharing agreements to align provider and customer incentives around a shared efficiency objective. Before the contract begins, the parties agree on a should-cost baseline. If the provider delivers below that baseline, savings are shared at a pre-agreed ratio—typically 50/50 for mature programs, sometimes 60/40 or 70/30 in favor of the party that assumed more risk (Defense Acquisition University [DAU], 2019; Schuster, 1993). Both parties are better off than under a pure cost-plus arrangement; both are aligned on the shared objective of driving cost below baseline; neither has an incentive to conceal efficiency opportunities from the other.

Commercial Gainsharing vs. Defense WCF

Commercial: $Provider\ Gain = (Should\ Cost - Actual\ Cost) \times 50\%$

$$Customer\ Gain\ =\ (Should\ Cost\ -\ Actual\ Cost)\ \times\ 50\%$$

$$Timing\ is\ usually\ a\ current\ contract\ year,\ thereby\ immediate, direct$$

DWCF: $Provider\ Gain\ =\ (Should\ Cost\ -\ Actual\ Cost)\ \times\ 0\%$

$$Customer\ Gain\ =\ Full\ savings, deferred\ 18-24\ months\ as\ AOR$$

The defense should-cost analysis is already mandatory under FAR 15.407-4. The provider has zero financial stake in beating should-cost. The customer receives the benefit two years later as rate adjustment. Therefore, neither party has a current-year financial incentive to drive efficiency.

The gainsharing formula converts the adversarial cost-plus dynamic into a cooperative one. The provider has a direct financial incentive to find efficiencies and share them with the customer. The customer has a direct financial incentive to provide the provider with the information, access, and flexibility needed to find those efficiencies. The should-cost baseline is the shared reference point against which both parties measure performance. The DWCF has no gainsharing mechanism despite having a mandatory should-cost requirement.

A Reinterpreted DWCF Is Needed: Innovation Working Capital Fund

The current Defense WCF, as examined in detail, is a cost-recovery instrument that operates with mechanical fidelity to its 1949 statutory purpose—and consequently generates break-even solvency, full-cost visibility, zero innovation capital, and, more importantly, is insufficient for the sustainment problem the Department faces. The commercial working capital fund, by contrast, operates a rate equation with structurally identical inputs but categorically different outputs: cost recovery, operating margin, reinvestment premium, and risk reserve, producing 3%–8% of revenue in self-financed innovation capital and 3-to-1 to 6-to-1 leverage against contracted revenue streams (BCG, 2024a; KPMG, 2023; Preqin, 2024).

This paper proposes a new DWCF model, the Innovation Working Capital Fund (IWCF), which draws on the benefits of a commercial WCF but still complies with the statutory constraints of 10 U.S.C. § 2208 and does not require a new appropriation or new legislative authorities. It preserves the aforementioned benefits of the DWCF while authorizing innovation, incentivizing a direct feedback loop, and crowding in in-kind private capital. It assigns the PAE the integrator of the WCF that can function like a GP to make portfolio allocation decisions, exercise governance rights, and set performance expectations. The ordering activities function as LPs, supplying capital through their sustainment appropriations in exchange for defined service delivery.

More than just an accounting function, the IWCF is modeled to realign incentives through the following mechanics:

- Streamlines the feedback loop between the warfighter and industry, which delivers the capability. In a post-production contract under the current requirements process, changes are made through a maintenance reporting system, which then has to undergo a contract modification through the contracting office, which is subject to the budget window and funding. The supplier has no in-year financial motive in the urgency of the solution.
- Promotes collaboration between the PAE-led IWCF and vendor consortia built on a codified reference architecture to lower barriers to entry, with structured financial equity and offsets against government equity.

- Bakes in an innovation premium in the prospective cost component in the rate, which self-finances innovation capital and eliminates reliance on additional appropriations or laborious contract modifications.
- As the integrator, the PAE improves the marketplace through public-private partnership structures that attract more competition and private capital. It can create a competitive refresh window that allows non-traditional firms to participate at MOSA-defined interfaces, providing ongoing market access that converts one-time demonstration projects into recurring revenue relationships.

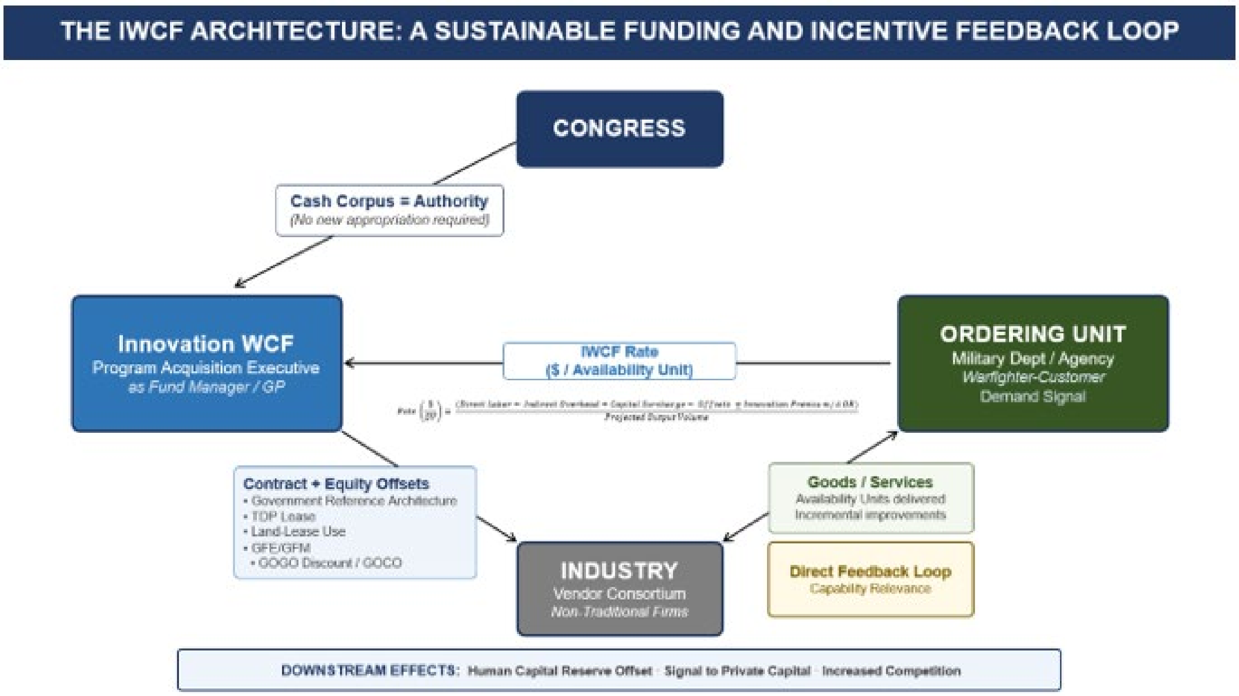

Note. The IWCF architecture as a sustainable funding and incentive feedback loop.

Congress → PAE IWCF
Cash Corpus = Authority. Congress provides § 2208 authority, new appropriation not required..

Ordering Unit → PAE IWCF
IWCF Rate paid per Availability Unit. Denominator shifts from Direct Labor Hours to system-mission-capable-days delivered.

PAE IWCF → Industry
Contract + Equity Offsets: Government Reference Architecture, TDP lease, GFE/GGM (GOGO Discount and GOCO Use)GOCO/GOGO, Land lease agreements

Industry → PAE IWCF
Gainshare Return: revenue share 4–8%, Reliability Bond 2%, surge capacity 20–30% callable, and affirmative data rights transfer.

Industry → Ordering Unit
Goods and services delivered as Availability Units, with incremental reliability improvements reflected in the rate.

Ordering Unit → Industry
Direct Feedback Loop on capability is streamlined between demand and supply side

Figure 2: The IWCF Architecture: A Sustainable Funding and Incentive Feedback Loop

The IWCF Rate Mechanics

The IWCF rate is still based on the DWCF within Financial Management Regulation compliance but incorporates the efficiencies of the commercial model by replacing the reinvestment premium with an innovation premium that is absorbed in the AOR within the rate period. It also hedges the capital surcharge with government equity offsets. The denominator is also now measured in output volume, defined as one system-mission-capable-day for a specified platform, and not labor hours, thus rewarding performance over consumption.

The IWCF is a single-rate structure decomposed into five components:

$$\boldsymbol{Rate}\left(\frac{\$}{\boldsymbol{OU}}\right) = \frac{(\boldsymbol{Direct\ Labor} + \boldsymbol{Indirect\ Overhead} + \boldsymbol{Capital\ Surcharge} - \boldsymbol{Offsets} \pm \boldsymbol{Innovation\ Premium/AOR})}{\boldsymbol{Projected\ Output\ Volume}}$$

The commercial cost floor was decomposed to the status quo DWCF components of direct labor and indirect overhead. Capital surcharge also remains the same—depreciation recovery on government-owned equipment—but it is offset by government equity offsets modeled after the aforementioned commercial gainsharing.

Government Equity Offsets

Four categories of government equity are eligible for gainsharing integration into the rate structure. Government equity can be currently unpriced or under-priced in the relationship with industry, but each can be converted into rate leverage through instruments that existing law already permits.

Technical Data Package Rights

Under Defense Federal Acquisition Regulation Supplement (DFARS) 252.227-7013, the government holds default unlimited data rights in items developed with exclusive government funding and in technical data pertaining to the form, fit, and function of items delivered under government contracts. In practice, the government owns substantial data rights by operation of law in many cases where commercial practice then permits the contractor to operate as though the government's rights did not exist—either through inadequate government TDP curation, through ambiguity about which specific data is subject to government rights, or through the contractor's practical monopoly on operational use of the data. The GAO (2025b) has documented that a majority of sole-source contract justifications cite inadequate technical data as the basis, even when the underlying legal entitlement of the government is unambiguous.

The government's unexercised TDP rights are the single most valuable unmonetized asset in the defense industrial relationship. Every dollar of sole-source premium the government pays on a sustainment contract is, in effect, a premium for the contractor's practical control of data the government legally owns. Converting that control into rate leverage—through contractual mechanisms that make TDP access operationally as well as legally available—is the highest-leverage gainsharing opportunity in the architecture.

Under the IWCF, the PAE maintains an annual sole-source premium benchmark by commodity category, documenting the implicit premium the government has historically paid for sustainment services where TDP access was limited. Where TDP access is structurally available, the premium the government would otherwise pay is avoided. The avoided premium is captured as the TDP Offset and deducted from the Cost Floor.

Real Property Offset: Outleasing and Enhanced Use Leases

The Department owns real property infrastructure—including five hard-iron Army depots (Anniston, Corpus Christi, Red River, Tobyhanna, and Watervliet), three arsenals, two munitions production plants, and multiple Air Force Air Logistics Complexes and Navy shipyards, all government-owned—valued in aggregate at well over $10 billion in replacement cost. Comparable commercial industrial capacity would cost a private operator approximately $18 to $25 per square foot per year to lease in current markets (CBRE, 2024; JLL, 2024). Across the government-owned defense industrial footprint, the government is effectively foregoing hundreds of millions of dollars annually in market-rate lease revenue that it could either collect or exchange for reduced rates on the services the facilities produce.

Under 10 U.S.C. § 2667, the Secretary of Defense is explicitly authorized to lease non-excess Department property, including at below-market rates where the public interest is served and in-kind consideration is accepted (CRS, n.d.). The proposed IWCF rate-structure mechanism captures the foregone lease value as reduced WCF cost to the government—converting the lease discount from an unpriced subsidy into an explicit rate offset.

Under the IWCF model, the PAE, as fund manager, commissions an annual independent valuation of market-rate lease equivalents for each GOCO and GOGO facility using standard commercial real estate methodologies. The difference between the market-rate lease equivalent and the actual lease rate (which may be zero at GOGO facilities and below-market at GOCO facilities) constitutes the Real Property Offset. The offset is deducted from the Cost Floor (Direct Labor + Indirect Overhead + Capital Surcharge) on a per-unit basis, reducing the rate the PAE charges ordering activities. The valuation operates under existing 10 U.S.C. § 2667 authority and standard commercial real estate appraisal practice (Appraisal Institute, 2024).

Government-Furnished Material and Equipment (Non-Land)

The same structural condition exists across GOCO facilities in the defense industrial complex. The contractor operates the equipment; the government owns it; the products manufactured using the equipment are sold back to the government at rates that do not discount for the government's contribution of the capital asset base.

DFARS 252.245-7001 and related Government-Furnished Property clauses establish the accountability framework for government-furnished equipment (GFE), and the clauses assume the contractor bears stewardship obligations for the equipment. What the clauses do not establish—and what the proposed gainsharing architecture introduces—is a valuation framework that recognizes the contractor's free use of the GFE as an implicit subsidy that should be priced into the rate.

Under the IWCF model, the PAE commissions an annual valuation of the GFE and tooling in active use at each facility, applying standard commercial depreciation and opportunity-cost methodologies (Equipment Leasing and Finance Association [ELFA], 2024). The annual opportunity cost of the GFE/M, or what a commercial operator would pay to lease equivalent equipment in a capital-asset leasing market, constitutes the GFE/M Offset. The offset is deducted from the Cost Floor. The PAE captures the value of its capital contribution; the contractor retains operational use of the equipment without separate lease payments.

Configuration Management and Process Data Generated on Government Equipment

A fourth category is emerging rapidly as digital engineering, automation, and technological improvements are required to meet the operational needs of the battlefield. When a contractor performs government-funded work using GFE or government-furnished material (GFM) in a government-owned facility, the process data generated during that work—the machine settings, tolerances, inspection results, and configuration records—is generated with government resources and pertains to government property. The default data rights framework under DFARS 252.227-7013 suggests this data belongs to the government, but few contracts assert the claim explicitly, and fewer still build the administrative infrastructure to capture, curate, and use the data (DoD, 2024b).

Under the proposed model, this data is treated as a government-contributed asset in the gainsharing relationship: The contractor retains operational use during the contract period, but the government retains the right to license the data to alternative providers during solicitation, to use the data for should-cost baseline validation, and to require the data be maintained in a digital thread format that supports the Department's configuration management and predictive maintenance objectives. Thus, the data rights belong to the government, but the administrative infrastructure to exercise them is converted into rate leverage.

Under the IWCF, where the contractor is required to deliver digital thread configuration data, process data, and predictive maintenance telemetry in a government-accessible format, the commercial value of that data—measured against comparable data licensing transactions in the commercial sectors—is captured as the Process Data Offset (BCG, 2024b; Deloitte, 2023). The offset is modest in size compared to the other three components but grows materially over time as digital engineering and technology become standard practice and as the commercial value of curated industrial data increases.

These government offsets are menu options that industry can elect into based on the use case, so the offset will vary. For instance, the Government Equity Offset captures:

- (a) Foregone lease value on GOCO/GOGO real property → 1.5%–2.5% of rate.
- (b) Amortized value of Government-Furnished Equipment → 1.0%–2.0% of rate.
- (c) Avoided sole-source premium from TDP access rights → 2.0%–4.0% of rate.
- (d) Process data licensing value (digital thread) → 0.5%–1.5% of rate.

Total Government Equity Offset: approximately 5%–10% of rate.

The structural effect of the offset is that the government pays a rate 5% to 10% below what the same services would cost if procured from a commercial provider with no government-contributed assets. Industry is no worse off, because the assets industry uses were not industry's capital in the first place. The government captures the value of what it already owns, in the form of a lower rate on the services it purchases back. Both parties have made the gainshare explicit, priced it against market benchmarks, and made the arrangement auditable.

A critical distinction of the IWCF's gainsharing is its architecture. The Government Equity Offset is structured as an explicit gainshare in which both parties are measurably better off than under the alternative (DAU, 2019; Schuster, 1993). Industry gives up market-rate lease payments on GOCO facilities, market-rate leasing income on GFE, sole-source premium revenue on TDP-protected work, and commercial data licensing revenue on process data—in exchange for reduced-margin contracts that nonetheless provide guaranteed multi-year revenue, reduced market-entry cost, access to government infrastructure at below-market effective rates, and a defined innovation return window. The government gives up the right to charge market rates for its assets—in exchange for the above example rate of 5% to 10% below commercial equivalent, a defined feedback loop to supplier performance, and structured private capital co-investment through the partnership architecture.

The arithmetic of the gainshare is favorable to both parties precisely because each party values what the other gives up differently. The government values reduced sustainment cost more than it values forgone lease revenue; the lease revenue the government would receive under a market-rate structure would flow to the Treasury as miscellaneous receipts, disconnected from the sustainment operation that created the value (Miscellaneous Receipts Statute, n.d.). Under the gainshare, the value flows directly to reduced rate on the sustainment services the PAE purchases. Industry values guaranteed multi-year demand and reduced market-entry cost more than it values marginal additional margin on facility lease payments it would otherwise receive and then deploy into general corporate operations. The structure makes both parties' value curves visible and aligns the arrangement accordingly.

Favorable Rates Through Structured Equity, Not Subsidy

The Government Equity Offset is distinct from a subsidy in a specific and important sense. A subsidy is a government payment that reduces a party's cost below what that party would otherwise bear. The Offset is a pricing recognition of assets the government has already contributed to the production relationship—assets whose cost has already been borne by the government through prior appropriations and whose value industry is already benefiting from on a free basis under the current structure. The Offset makes the existing implicit subsidy explicit and captures its value as rate leverage rather than leaving it unrecognized in the relationship.

This distinction matters to audit defensibility. An OSW Comptroller or GAO reviewer examining the rate will see each component of the Government Equity Offset tied to a market-benchmark valuation, documented through an independent third-party appraisal, and applied as a mathematical deduction from the Cost Floor. The reviewer can trace each dollar of rate reduction to a specific government-owned asset and to the market-rate equivalent that would otherwise have been charged. Therefore, the audit trail is complete: Valuations are independent, the methodology is standard commercial practice, and the final rate is transparently reconcilable to both the Cost Floor baseline and the Offset components.

Innovation Premium Under Existing Authority

Section 2208(a) limits WCF authorization to finance inventories of such supplies as the Secretary may designate and provide working capital for such industrial-type activities as the Secretary may designate. The chartered scope is industrial-type activities and supply inventory management. It does not extend to independent research and development, venture investments, equity stakes in suppliers, or procurement of major end items. The Innovation Premium can only finance activities that fall within the chartered scope—which is a material constraint on what kinds of innovation the premium can support, but also an affirmative definition of what kinds of innovation it can (Working-Capital Funds, n.d.).

However, the FY2024 NDAA amendment to § 2208(c) explicitly added to the list of WCF-chargeable costs: the cost of the procurement and qualification of technology-enhanced maintenance capabilities that improve either reliability, maintainability, sustainability, or supportability and have, at a minimum, been demonstrated to be functional in an actual system application or operational environment (NDAA, 2023). Furthermore, the FY2026 NDAA codification of the PAE authority and the Warfighting Acquisition System establishes Congressional intent that the Department manage sustainment at the portfolio level and modernize the industrial base through executive rate-structure management (NDAA, 2025).

The IWCF Innovation Premium operates in the same legal category as the Capital Surcharge—a recognized prospective cost-recovery component with an established precedent. As authorized under 10 U.S.C. § 2208(e)(1), the Capital Surcharge recovers the cost of capital investments the fund makes during the rate period. In parallel, the Innovation Premium recovers the cost of technology-enhanced maintenance activities the fund performs during the rate period.

The Innovation Premium can finance five qualifying activity categories, each of which falls within the WCF's chartered scope under § 2208(a) and each of which incurs its cost during the rate period.

Table 3: Innovation Premium Qualifying Activity Categories Under Existing Working Capital Fund Authority

(DoD, n.d.-c, n.d.-d, 2024; Working-Capital Funds, n.d.)

Qualifying Category	**Activity**	**WCF Authority Precedent**
Technical Data Package Acquisition	One-time payment to original manufacturer or prior developer to acquire or expand government rights in technical data required for sustainment operations.	10 U.S.C. § 2208(c) technology-enhanced maintenance language.
Alternative Supplier Qualification	Payments to testing laboratories, certification bodies, and qualification engineers to qualify second-source suppliers for parts currently sole-sourced.	DLA Supply Chain Management activity group precedent.
Digital Engineering Infrastructure	Software licenses, sensor installations, and data platform configuration to support configuration management and predictive maintenance on platforms the fund sustains.	DISA Information Services activity group precedent; DoD Instruction 5000.97.
Condition-Based Maintenance Tooling	Diagnostic equipment, test stands, and reliability-monitoring systems placed in service at depot locations during the rate period.	Existing Capital Investment Program under DoD FMR Vol. 11B Ch. 6.
Competitive Evaluation Activities	Costs of soliciting, evaluating, and qualifying alternative providers during the Competitive Refresh Window.	10 U.S.C. § 2208(e)(1) rate-setting authority for evaluation activities.
All five categories	Each produces cost incurred in the rate period; each has direct precedent in existing WCF practice.	Combined authority basis is affirmative and established.

The Innovation Premium is a cost component of the rate itself, charged prospectively for innovative activities the fund will perform in the rate period. The PAE, as fund manager, publishes the annual rate with an example Innovation Premium line of 3% to 5% of the rate, designated for specific, PAE-verified innovation investments planned for that year. Ordering activities pay the premium at the time of billing. The fund spends it on the designated purposes during the year. Any unspent balance at year-end is returned through the AOR mechanism in the subsequent rate cycle, preserving the break-even mandate over the operating cycle. These funds must be spent or returned via AOR within the rate period, confirming no multi-year retention.

This is authorized under 10 U.S.C. § 2208(e)(1) because it is how the DoD Financial Management Regulation already treats the Capital Surcharge component—the surcharge funds capital investments made during the rate period, not retained for future speculative use. The Innovation Premium operates on the identical legal theory that distinguishes an authorized rate component from an unauthorized fund balance accumulation (DoD, n.d.-c).

Justification Under the Expanded Capital Investment Program Scope

The Capital Investment Program, authorized under DoD FMR Volume 11B Chapter 6, already permits the WCF to fund capital improvements with expected useful lives beyond one year. Expanding the CIP definition to include specific categories of innovation investment treats these investments as capital assets of the fund rather than as retained surplus.

The investments would appear on the fund's balance sheet, then be depreciated over their useful lives (typically 5 to 10 years depending on asset category), and recovered through the Capital Surcharge in future rates. In essence, the fund buys assets that produce an operational benefit during their useful lives, which reduces the fund's cost of operations in each depreciation period. These industrial-type activities financed on a revolving basis with the qualifying innovation categories can be explicitly recognized as capital assets. However, this scope expansion may require a policy memorandum. The scope expansion depends on the Comptroller's interpretative authority and may require a policy memorandum at the Under Secretary level, but not a statutory amendment.

Furthermore, Section 2208(l) authorizes the WCF to maintain a stabilization reserve against demand volatility and price fluctuation. The statute requires the reserve be held against future working capital requirements, and it does not restrict the form in which the reserve is held. Under this interpretation, the PAE can direct the stabilization reserve's holding into innovation investments that simultaneously qualify as working capital—specifically, prepaid supplier qualification (reducing future qualification cost requirements), pre-positioned alternative-source inventory (reducing future supply-chain disruption working capital), and efficiency-improving digital infrastructure (reducing future inventory carrying costs). In other words, they are the working capital itself, deployed into assets that reduce the fund's future working capital requirements while remaining available as a reserve against volatility (Working-Capital Funds, n.d.).

When the alternative supplier is qualified, the fund's future working capital requirement is reduced because sole-source risk premiums in inventory and pricing decline. When the digital engineering infrastructure is built, inventory carrying costs decline because demand forecasting improves. When the pre-positioned inventory is in place, the working capital requirement against supply-chain disruption declines. The reserve is being held against future needs but held in a form that actively reduces those future needs.

A similar example of this expanded mechanism can be extrapolated from 10 U.S.C. § 4123, which authorizes the director of a defense laboratory or test organization to charge customer activities a fixed fee that may not exceed 4% of costs to fund infrastructure projects.

The PAE-Integrated Vendor Consortium

This model also reinterprets the PAE's relationship with its vendor base from a series of bilateral program-level contracts into an integrated consortium operating under a single portfolio-level governance framework. The PAE-Integrated Vendor Consortium, structured through existing OTA authority under 10 U.S.C. § 4022, creates a standing vehicle through which alternative suppliers can be qualified and deployed without requiring bespoke program-level solicitations. Under this model, it would also manage and enforce a Government Reference Architecture (GRA) compliance standard to ensure interoperability and technical consistency across vendors and only source firms that opt into GRA compliance. The Air Force has already proven this example through its Collaborative Combat Aircraft (CCA) program, in which the program integrator's role is to enforce that all vendor candidates must be compliant with the Government Avionics Reference Architecture (GARA) and compete to deliver modular, interoperable AI services rather than bespoke systems (U.S. Air Force, 2024).

The consortium structure produces three effects the current fragmented contracting environment cannot. First, it reduces the transaction cost of bringing non-traditional firms into defense sustainment, because the consortium provides a standing vehicle rather than requiring every new entrant to navigate individual program solicitations. Second, it enables portfolio-level cost-reduction initiatives that span multiple programs simultaneously because the consortium's governance structure can modify obligations across all members concurrently. Third, it creates a recurring venue for competitive pressure on incumbents.

Downstream Effects of the IWCF Model

The Signal to Private Capital

Private capital deployment into defense sustainment markets has historically been limited because the revenue commitments the government could offer were annual, appropriation-dependent, and subject to continuing resolution volatility. The IWCF architecture, operated by a PAE as fund manager with multi-year rate stability and outputs anchored to the fund's revolving working capital, creates the counterparty profile that commercial infrastructure funds recognize (Deloitte, 2022; Preqin, 2024). The effect is observable in capital-expenditure runway: A non-traditional firm considering entry into defense sustainment faces materially shorter time-to-revenue when the counterparty is a PAE managing a multi-year demand certificate than when the counterparty is a program office managing an annual contract modification.

Increased Competitive Landscape and Human Capital Reserve Credits

The proposed architecture extends two downstream benefits. First, small and mid-sized companies gain immediate access to government revenue through the PAE-Integrated Vendor Consortium, compressing engagement-to-first-revenue timelines from the current 18-to-36-month traditional contracting cycle to roughly 3–9 months under new contract strategies. Second, the architecture opens the policy pathway for a Reservist Industrial Skills framework connecting civilian industrial work on qualifying contracts to the reservist's military career through two complementary instruments: retirement points credit for qualifying civilian industrial work, implementable under the Secretary of Military Department's existing authority under 10 U.S.C. Chapter 1209 to define qualifying activities, and a Reservist Industrial Skills Tax Credit extending the established I.R.C. § 45P framework from the Heroes Earnings Assistance and Relief Tax Act of 2008, under which employers whose reservists perform qualifying work on contracts issued through the fund receive a tax credit calibrated to wages paid for that work. This also gives the Department a pre-identified surge workforce—aligning contract access, workforce development, and industrial base resilience.

Conclusion

A reimagined capability-based WCF offers not only a funding mechanism to 'innovate sustainment' but also market-aligned feedback loop that streamlines demand, promotes collaboration between the PAE and vendor consortia, maximizes industry partnerships and incentives through offsets and public-private partnerships, and ensures that our warfighters can meet the needs of today's battlefield, and not that of a prior budget cycle.

References

Anti-Deficiency Act, 31 U.S.C. § 1341 (n.d.). https://www.law.cornell.edu/uscode/text/31/1341

The Appraisal Foundation. (2024). *Uniform standards of professional appraisal practice (USPAP)* (2024 ed.). Appraisal Standards Board. https://appraisalfoundation.org/pages/uspap

Assistant Secretary of the Army (Financial Management & Comptroller). (2025, June). *Army working capital fund FY2026 budget estimates*. Department of the Army. https://www.asafm.army.mil/Budget-Materials/

Blackstone Inc. (2025). *Annual report (Form 10-K) for fiscal year 2024*. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/0001393818/000119312525042469/d912273d10k.htm

Boston Consulting Group. (2024). *Where's the value in AI?* https://www.bcg.com/publications/2024/wheres-value-in-ai

Boston Consulting Group & World Economic Forum. (2021, January). *Data excellence: Transforming manufacturing and supply systems*. https://www.bcg.com/about/partner-ecosystem/world-economic-forum/data-value-manufacturing

CBRE. (2024). *U.S. real estate market outlook 2024: Industrial & logistics*. https://www.cbre.com/insights/books/us-real-estate-market-outlook-2024/industrial

CFA Institute. (2024). *Working capital management: CFA program level I curriculum, corporate issuers module*.

Congressional Research Service. (n.d.). *Military installations, real property, and land management* (CRS Report No. IF11309). https://www.congress.gov/crs_external_products/IF/PDF/IF11309/IF11309.2.pdf

Congressional Research Service. (2024a, May). *PPBE reform commission final report recommendations: Issues for Congress* (CRS Report No. IN12372). https://crsreports.congress.gov/product/pdf/IN/IN12372

Congressional Research Service. (2024b, December). *Defense primer: Defense working capital funds* (CRS Report No. IF11233). https://crsreports.congress.gov/product/pdf/IF/IF11233

Defense Acquisition University. (2019). *Gainsharing guide*. https://www.dau.edu/

Defense Laboratory and Test Organization Fees, 10 U.S.C. § 4123 (n.d.). https://www.law.cornell.edu/uscode/text/10/4123

DefenseScoop. (2025, November 14). *Army introduces sweeping reform of its acquisition structure*. https://defensescoop.com/2025/11/14/army-acquisition-reform-2025/

Deloitte. (2022). *Infrastructure finance: The private capital opportunity*. https://www2.deloitte.com/

Deloitte. (2023). *Industrial data monetization*. https://www2.deloitte.com/

DFARS 216.405-1, Cost-Plus-Incentive-Fee Contracts (n.d.). https://www.acquisition.gov/dfars/216.405-1-cost-plus-incentive-fee-contracts

DFARS 252.227-7013, Rights in Technical Data—Noncommercial Items (n.d.). https://www.acquisition.gov/dfars/252.227-7013-rights-technical-data-noncommercial-items

DFARS 252.227-7014, Rights in Noncommercial Computer Software and Noncommercial Computer Software Documentation (n.d.). https://www.acquisition.gov/dfars/252.227-7014-rights-noncommercial-computer-software-and-noncommercial-computer-software-documentation

DFARS 252.227-7019, Validation of Asserted Restrictions—Computer Software (n.d.). https://www.acquisition.gov/dfars/252.227-7019-validation-asserted-restrictions-computer-software.

DFARS 252.245-7001–252.245-7004, Government-Furnished Property Clauses (n.d.). https://www.acquisition.gov/dfars/part-252-clauses

DoD. (n.d.-a). *Accelerate the Procurement and Fielding of Innovative Technologies (APFIT) program overview*. https://ac.cto.mil/apfit/

DoD. (n.d.-b). *DoD financial management regulation* (DoD 7000.14-R, Vol. 3, Chapter 15). https://comptroller.defense.gov/FMR/

DoD. (n.d.-c). *DoD financial management regulation* (DoD 7000.14-R, Vol. 11B, Chapter 1, §§ 010201–010207). https://comptroller.defense.gov/FMR/vol11b_chapters/

DoD. (n.d.-d). *DoD financial management regulation* (DoD 7000.14-R, Vol. 11B, Chapter 6). https://comptroller.defense.gov/FMR/vol11b_chapters/

DoD. (2022, February 15). *State of competition within the defense industrial base*. https://media.defense.gov/2022/Feb/15/2002939087/-1/-1/1/state-of-competition-within-the-defense-industrial-base.pdf

DoD. (2024, February). *Digital engineering* (DoD Instruction 5000.97). https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500097p.PDF

DoD. (2025, November 10). *Acquisition transformation strategy: Rebuilding the arsenal of freedom*. https://www.war.gov/

DoD, Office of the Under Secretary of Defense (Comptroller). (2025a, March). *Department of Defense FY2026 budget overview*. https://comptroller.defense.gov/Budget-Materials/

DoD, Office of the Under Secretary of Defense (Comptroller). (2025b, June). *Defense-Wide Working Capital Fund FY2026 budget estimates*. https://comptroller.defense.gov/Budget-Materials/

Duffey, M. (2026, March). *Testimony before Congress* [Under Secretary of War for Acquisition and Sustainment].

Employer Wage Credit for Active Duty Members of the Uniformed Services (Heroes Earnings Assistance and Relief Tax Act of 2008), 26 U.S.C. § 45P (n.d.). https://www.law.cornell.edu/uscode/text/26/45P

Equipment Leasing and Finance Association. (2024). *State of the equipment finance industry report*. https://www.elfaonline.org/

Exec. Order No. 14265, 90 Fed. Reg. 15,621 (2025, April 9). https://www.federalregister.gov/documents/2025/04/14/2025-06347/modernizing-defense-acquisitions-and-spurring-innovation-in-the-defense-industrial-base

FAR 15.407-4, Should-Cost Review (n.d.). https://www.acquisition.gov/far/15.407-4

FAR 16.405-1, Cost-Plus-Incentive-Fee Contracts (n.d.). https://www.acquisition.gov/far/16.405-1

FAR Part 48, Value Engineering (n.d.). https://www.acquisition.gov/far/part-48

Federal News Network. (2025, October 29). *A new framework aims to build a bridge across the defense acquisition valley of death*. https://federalnewsnetwork.com/

GAO. (2023). *F-35 sustainment: DOD needs to address challenges affecting readiness and cost transparency* (GAO-23-105262). https://www.gao.gov/products/gao-23-105262

GAO. (2024, February 29). *Weapon system sustainment: DOD identified operating and support cost growth but needs to improve the consistency and completeness of information to Congress* (GAO-24-107378). https://www.gao.gov/products/gao-24-107378

GAO. (2025a). *Weapon system sustainment: DOD can improve planning and management of data rights* (GAO-25-107468). https://www.gao.gov/products/gao-25-107468

GAO. (2025b, June 11). *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

JLL. (2024). *U.S. industrial market statistics & trends*. https://www.jll.com/en-us/insights/market-dynamics/industrial-market-statistics-trends

KPMG. (2023). *Global managed services industry survey*. https://kpmg.com/

Lease of Non-Excess Property of Military Departments and Defense Agencies, 10 U.S.C. § 2667 (n.d.). https://www.law.cornell.edu/uscode/text/10/2667

Macquarie Asset Management. (2024). *Infrastructure fund series disclosures*. https://www.macquarieam.com/

Miscellaneous Receipts Statute, 31 U.S.C. § 3302 (n.d.). https://www.law.cornell.edu/uscode/text/31/3302

MITRE. (2025). *Transition maturity framework* [Developed with OSD Acquisition and Sustainment/Operational Energy Innovation Organization]. https://www.mitre.org/

National Defense Authorization Act for Fiscal Year 2024, Pub. L. No. 118–31, 137 Stat. 136 (2023). https://www.congress.gov/bill/118th-congress/house-bill/2670

National Defense Authorization Act for Fiscal Year 2025, Pub. L. No. 118–159 (2024). https://www.congress.gov/bill/118th-congress/house-bill/5009

National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119–__, §§ 1801–1802 (2025). https://www.congress.gov/

National Defense Industrial Association. (2023). *Performance-based logistics policy paper*. https://www.ndia.org/

National Security Act Amendments of 1949, Pub. L. No. 81–216, 63 Stat. 578 (1949). https://www.govinfo.gov/app/details/STATUTE-63/STATUTE-63-Pg578

Office of Strategic Capital. (2024). *FY2025 investment strategy*. DoD. https://www.cto.mil/osc/

Other Transactions for Prototypes, 10 U.S.C. § 4022 (n.d.). https://www.law.cornell.edu/uscode/text/10/4022

Other Transactions for Research, 10 U.S.C. § 4023 (n.d.). https://www.law.cornell.edu/uscode/text/10/4023

PPBE Reform Commission. (2024, March). *Final report of the Commission on Planning, Programming, Budgeting, and Execution Reform*. https://ppbereform.senate.gov/

Preqin. (2024). *Global infrastructure report 2024*. https://www.preqin.com/

Prohibition on Contracting for Performance of Depot-Level Maintenance and Repair, 10 U.S.C. § 2466 (n.d.). https://www.law.cornell.edu/uscode/text/10/2466

Purpose Statute, 31 U.S.C. § 1301 (n.d.). https://www.law.cornell.edu/uscode/text/31/1301

Reprogramming of Funds, 10 U.S.C. § 7013 (n.d.). https://www.law.cornell.edu/uscode/text/10/7013

Roque, A. (2025, December 19). *From Army contracting pause to Pentagon acquisition overhaul: 2025 review*. Breaking Defense. https://breakingdefense.com/2025/12/from-army-contracting-pause-to-pentagon-acquisition-overhaul-2025-review/

Sale of Ammunition, Munitions, and Components Thereof, 10 U.S.C. § 4543 (n.d.). https://www.law.cornell.edu/uscode/text/10/4543

SC World. (2026, April). *Acquisition reform is materializing, but the harder test still lies ahead*. https://scworld.com/

Schuster, M. H. (1993, September–October). Gainsharing: Sharing productivity with employees. *Harvard Business Review*. https://hbr.org/

Secretary of War. (2025, November 10). *Transforming the Defense Acquisition System into the Warfighting Acquisition System* [Memorandum]. DoW. https://www.war.gov/

U.S. Air Force. (n.d.). *Collaborative Combat Aircraft (CCA) program overview*. https://www.af.mil/

Working-Capital Funds, 10 U.S. Code § 2208 (n.d.). https://www.law.cornell.edu/uscode/text/10/2208

Quantifying Return on Investment in Digital Technical Data Packages for Combat Systems: A Life Cycle Perspective

John G. Fiore—is Vice President of Strategic Growth at Kitty Hawk Technologies in Dahlgren, VA, where he leads innovation initiatives, identifies emerging opportunities, and builds mission focused partnerships that strengthen warfighting capability and community impact. A former member of the Senior Executive Service, he previously served as the Director for Acquisition and Sustainment at the Missile Defense Agency, overseeing enterprise wide acquisition and sustainment strategies. Before joining MDA, he was the Technical Director of the Naval Surface Warfare Center, Dahlgren Division, guiding a large technical workforce in delivering advanced capabilities to the Fleet.

Glenna Miller—is a premier technical talent who brings threat engineering to life in ways vital to national security. She has more than 35 years of experience across DoD technical, engineering, program management, and leadership roles, serving within government, the military, and as a contractor supporting Navy and Air Force programs. Her background includes serving as Deputy Threat Engineering Division and Threat Engineering Branch Head at NSWC/DD and leading the PEO IWS threat working group. Her expertise spans agile and traditional systems engineering, MBSE, S&TI analysis, and Surface Navy Warfare requirements. She now serves as Kitty Hawk's technical program manager.

Chris Petrin—is a systems engineer specializing in C UAS, remote sensing, and the application of digital engineering principles to accelerate effective designs and warfighter capabilities. At Kitty Hawk Technologies, he develops physics based and statistical flight models, advanced RCS characterization methods, and modular simulation architectures. Previously, at the Oklahoma Aerospace Institute for Research and Education, he led interdisciplinary teams advancing UAS/C UAS test and evaluation for DoD and DHS sponsors. He has designed RF measurement platforms, executed major C UAS test campaigns, and synthesized intelligence to shape test vignettes. Chris holds BS and MS degrees and is completing a PhD in Mechanical and Aerospace Engineering.

Abstract

Combat system acquisition is shifting toward digital engineering (DE) as a strategic necessity for accelerating warfighting capability. Traditional Technical Data Package (TDP) development relies on physical integration and live testing, often causing delays and high costs. Digital TDPs instead enable virtual performance exploration, executable mission threads, and early configuration control. Even in sustainment phases, they support faster updates and reduce risk through model-based validation.

Using ACAT I benchmarks, this paper quantifies DE return on investment (ROI) across the life cycle. One case draws on early Model-Based Test and Evaluation (MBT&E) applied to the AEGIS Ballistic Missile Defense baseline, reducing live-fire tests by 50%, avoiding about $222 million in costs, and producing a reusable modeling and simulation environment. A second case examines the Threat Digital Twin Advanced Technology Demonstration, showing how validated digital twins improve design performance, reduce test burden, and enable reuse across variants.

Results show DE delivers ROI at any life cycle stage, with the largest gains in Operations and Support, where most life cycle costs occur. Key benefits include fewer live tests, faster integration and validation, improved configuration management, and greater stakeholder confidence. DE transforms acquisition into a faster, more resilient force multiplier.

Introduction

Major Defense Acquisition Programs (MDAPs) operate under extraordinary technical complexity, decades-long service lives, and statutory oversight that magnify the consequences of engineering decisions made early in the life cycle. Yet despite more than a decade of policy emphasis on digital engineering (DE), the Department still lacks a quantitative, evidence-based understanding of where, when, and how DE produces measurable return on investment. This absence of analytic grounding is especially problematic for program managers, who are rarely provided discretionary resources to adopt new engineering practices or toolchains, particularly in legacy programs where technical debt is highest, documentation is fragmented, and budgets are most constrained. In this environment, DE is often mandated as a compliance requirement rather than justified as a beneficial life cycle investment, leaving programs without the evidence needed to defend or prioritize adoption. Policy has outpaced practice, creating a disconnect in which programs are expected to implement DE without the analytic basis needed to justify investment or the institutional support required to succeed.

The prevailing narrative in both policy and industry suggests that DE only produces meaningful return on investment (ROI) when adopted early, during Materiel Solution Analysis (MSA) or Technology Maturation and Risk Reduction (TMRR). While early adoption unquestionably provides the cleanest opportunity to establish a seamless digital thread, this assumption is rarely tested and often untrue. Many ACAT I programs begin DE adoption mid-stream or even during production and sustainment, yet still achieve measurable cost avoidance, risk reduction, and modernization agility.

This paper addresses that gap by examining ROI across two complementary scenarios. The first, and primary empirical case, draws directly on the author's experience as the former Deputy Technical Director and Chief Engineer for AEGIS BMD during the baseline development period (circa 2007–2014). It demonstrates how model-based test and evaluation principles reduced physical test volume, cost, and operational risk by more than $222 million, before such approaches were formalized in Department of Defense (DoD) policy. The second scenario examines the Threat Digital Twin ATD, a current program supported by Kitty Hawk Technologies, as a forward application of those proven principles, showing how the structural conditions for measurable life cycle ROI have been deliberately established in a mission area where no adequate physical test or training asset previously existed. Together, these scenarios provide both retrospective evidence and prospective framework demonstrating that DE generates significant ROI even when adopted mid-stream or late, and that the largest returns often emerge in O&S, where 60% to 70% of life cycle cost resides (GAO, 2024a).

Beyond illustrating ROI, this paper offers a practical framework for how programs can collect, curate, and structure the evidence needed to justify DE investment. By grounding the analysis in real program experience and aligning it with current DoD policy, the paper provides a defensible, actionable foundation for program managers seeking to adopt DE in environments where resources are limited, legacy constraints are significant, and the need for modernization is urgent.

Background

Defense Acquisition Context for ACAT I Programs

MDAPs, particularly those designated as ACAT I, operate under a uniquely demanding combination of technical complexity, statutory oversight, and decades-long service life. These programs progress through the DoD acquisition framework—MSA, TMRR, Engineering and Manufacturing Development (EMD), Production and Deployment (P&D), and Operations and Support (O&S)—yet the majority of life cycle cost and risk emerges only after Milestone C (entrance to the P&D phase). Across MDAPs, O&S consistently accounts for up to 70% of total life cycle cost, as documented in Selected Acquisition Reports (DoD, 2025). This dominance is driven by recurring maintenance and depot activity, continuous software sustainment and cybersecurity updates, parts obsolescence and diminishing manufacturing sources and material shortages (DMSMS) pressures, block upgrades required to pace evolving threats, and configuration drift across fielded units (DoD, 2025).

Traditional document-centric Technical Data Packages (TDPs) struggle to support these long-term demands; fragmented interface specifications, inconsistent configuration baselines, and limited traceability create brittleness that compounds over time. As a result, ACAT I programs face a structural challenge: the phases where most cost and risk accumulate are precisely the phases where traditional engineering artifacts are least effective. DE and model-based systems engineering (MBSE) directly target this gap by creating authoritative, persistent, and executable representations of the system—assets that remain useful throughout the life cycle rather than only during development.

Current DoD DE Strategy and Policy Environment

The DoD Digital Engineering Strategy (2018) remains the central policy anchor for DE adoption, articulating five goals: formalizing the development and use of models, establishing authoritative sources of truth, incorporating technological innovation, building supporting infrastructure, and transforming the culture and workforce (DoD, 2018). Importantly, these goals extend beyond early design and explicitly include sustainment, modernization, and test. Complementary policy reinforces this life cycle orientation, including DoDI 5000.97 on Digital Engineering for Defense Acquisition (DoD, 2023), DoDI 5000.61 on M&S Verification, Validation, and Accreditation (DoD, 2022), and 5000-series T&E policy permitting accredited M&S to replace or supplement physical testing like DoDI 5000.89 (DoD, 2020). A recent report from RAND further strengthens this policy support by noting that DE cost-benefit assessment is feasible at any life cycle stage and that systems-thinking and goal-based systems engineering are essential for realizing measurable returns (Whitehead et al., 2024).

These policies are aligned with larger trends across industries. The U.S. Government Accountability Office (GAO) surveyed 14 leading companies in product development and found extensive use of iterative design using modeling and simulation to create and maintain a digital thread along with validation using combinations of physical and digital prototypes (GAO, 2023). Digital twins, as an example of a key component of DE, are finding applications in manufacturing, aerospace, construction, agriculture, healthcare, retail, and mining (Attaran & Celik, 2023).

Together, these policy and analytic foundations support the central thesis of this paper: DE is a life cycle investment with quantifiable returns in sustainment, modernization, and test, not merely a front-end documentation replacement.

The Economics of DE Adoption Timing

A key finding from the RAND 2024 framework is that DE benefits are not limited to early adopters (Whitehead et al., 2024). While programs that begin DE during MSA or TMRR achieve the most seamless digital thread, meaningful ROI remains available at SRR, PDR, CDR, or even during production and O&S. Early adoption offers the highest theoretical return by enabling a full digital thread from concept and reducing integration risk entering EMD. Mid-stream adoption is the most common in real programs (GAO, 2025), and typically involves a digital retrofit of existing artifacts and yields improvements in integration, verification, and configuration control. Even late adoption during P&D or O&S may remain valuable when major block upgrades are planned, decades of service life remain, or obsolescence and software modernization pressures are high.

The economic logic is straightforward: when O&S represents 60% to 70% of life cycle cost, even modest efficiency gains can dwarf the cost of DE adoption. To illustrate this effect, Figure 1 shows the cumulative cost for the Illustrative System Life Cycle from the DoD O&S Cost-Estimating Guide (DoD, 2025). Two implementation cases are shown as well. In Case 1, an additional hypothetical 20% cost to implement DE is incurred during RDT&E, with variable percent cost savings in O&S. In Case 2, the hypothetical cost is incurred later to represent implementation during O&S, with benefits realized at the completion of the implementation. In Case 1 using the Illustrative System Life Cycle, only a 3.3% reduction in annual O&S costs is necessary to break even on the DE implementation investment. For Case 2, the break-even point becomes a 5.5% reduction in O&S costs. Both cases assume a worst case, that DE implementation yields no benefits until and only in O&S. Yet in both cases, modest O&S cost reductions can justify the implementation.

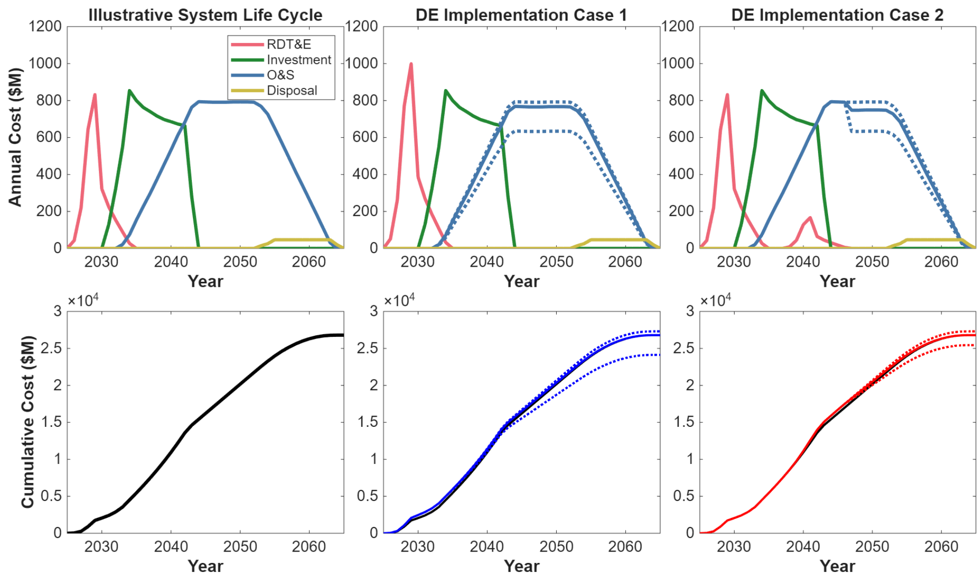

Figure 29. (Top) Annual Costs for the Illustrative System Life Cycle from DoD O&S Cost-Estimating Guide (DoD, 2025), and Two DE Implementation Cases. (Bottom) The Cumulative Cost Incurred for the Annual Costs Shown Directly Above. The dashed bands indicate 0% O&S savings (top dashed curve) and up to 20% O&S cost reduction (bottom dashed curve).

Figure 1 establishes the break-even threshold using DoD’s own published life cycle cost curve—a conservative, single-source test of the economic logic. Figure 2 broadens that analysis into a phase-by-phase ROM estimate synthesized across four policy and analytic sources, showing where the acquisition life cycle DE value accumulates and at what order of magnitude.

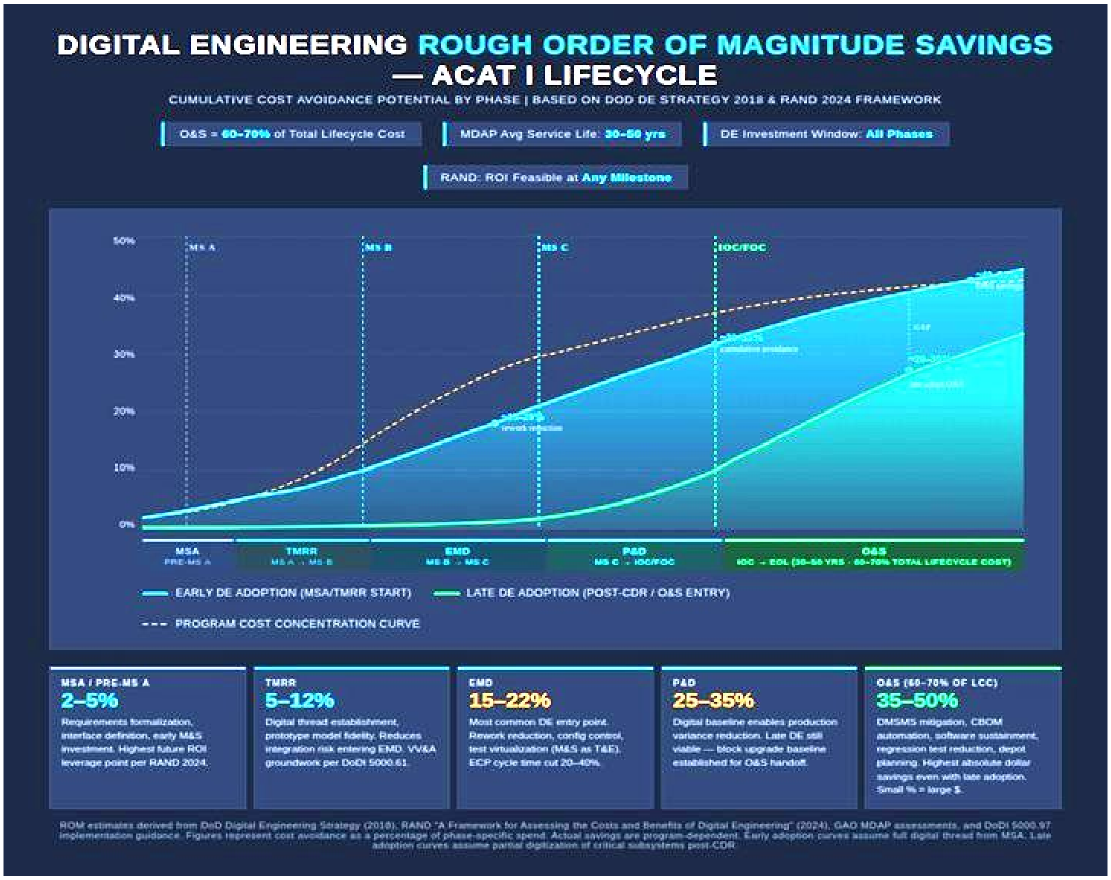

Graphic developed with AI assistance (Claude, Anthropic, 2025) based on the DoD Digital Engineering Strategy (2018); RAND Corporation, "A Framework for Assessing the Costs and Benefits of Digital Engineering" (2024); GAO MDAP program assessments; and DoDI 5000.97. Cost-avoidance values are rough-order-of-magnitude estimates and are program-dependent.

Figure 2. DE Rough Order of Magnitude Savings—ACAT I Life Cycle

Figure 2 illustrates cumulative cost-avoidance potential across the ACAT I life cycle. Early DE adoption beginning at MSA can yield 40%–45% cumulative avoidance by full operational capability (FOC)—roughly 15 percentage points greater than late adoption at CDR or O&S—but meaningful returns are available at every phase. Even late adoption produces 28%–35% cumulative avoidance, a compelling return given the absolute dollar scale of O&S. These ranges are rough-order-of-magnitude estimates derived from the DoD DE Strategy (DoD, 2018), RAND's 2024 framework (Whitehead et al., 2024), GAO MDAP assessments (GAO, 2025), and DoDI 5000.97 (DoD, 2023), and represent potential for cumulative, not annual, cost avoidance.

Table 1. DE ROM Cost Avoidance by Acquisition Phase (ACAT I Programs)

Acquisition Phase	ROM Cost Avoidance	Primary DE/MBSE Value Drivers	Why Savings Are Anticipated
MSA / Pre-MS A	2%–5%	Requirements formalization; interface definition; early M&S investment	Early clarity reduces downstream rework; highest leverage point per RAND 2024
Technology Maturation and Risk Reduction (TMRR)	5%–12%	Digital thread establishment; prototype model fidelity; VV&A groundwork	Reduces integration risk entering EMD; anchors early model credibility
Engineering and Manufacturing Development (EMD / MS B)	15%–22%	Rework reduction; configuration control; test virtualization; M&S as T&E	Cuts ECP cycle time by 20%–40%; reduces physical test load
Production & Deployment (MS C)	25%–35%	Digital baseline for production; variance reduction; upgrade pathway establishment	Late DE still viable; stabilizes configuration for O&S handoff
Operations & Support (IOC/FOC)	35%–50%	DMSMS mitigation; configuration-controlled parts automation; software sustainment; regression test reduction; depot planning	Highest absolute dollar savings because O&S is 60–70% of life cycle cost

Source: Fiore, J. G., "Digital Engineering ROM Cost Avoidance Methodology," internal working paper, 2025, derived from DoD DE Strategy (2018), RAND (2024), GAO MDAP assessments, and DoDI 5000.97.

Taken together, the acquisition realities, policy foundations, and economic analysis presented in this section establish a compelling case for DE as a life cycle imperative. ACAT I programs face an environment where 60% to 70% of total cost accumulates in O&S, where configuration drift and obsolescence pressures intensify over decades, and where traditional document-centric artifacts cannot keep pace with continuous software sustainment, cybersecurity, and block upgrades. The scenarios that follow demonstrate through concrete evidence how these principles translate into real reductions in test cost, operational risk, and life cycle burden.

Scenario 1—MBT&E: AEGIS BMD Baseline

(Personal Technical and Programmatic Reflection, John G. Fiore)

As a senior technical leader during the AEGIS BMD baseline development period (circa 2007–2014), MBT&E principles offered a powerful means of reducing cost, schedule pressure, and operational risk in a highly complex ACAT I environment. Although the formal digital-engineering constructs we reference today, such as MBTEMP, iTEMS, authoritative digital baselines, and enterprise digital threads did not exist circa 2008, the engineering team applied many of the underlying principles well before they were formally named or institutionalized.

At that time, Developmental Test, Operational Test, and Live Fire Test and Evaluation imposed significant burdens on the program, particularly because the AEGIS BMD kill chain was so tightly coupled that a failure in any segment invalidated the entire event. The Navy's Operational Test Agency (OPTEVFOR) initially requested 12 live-fire flight tests to demonstrate baseline performance, driven by the variability of threat environments, radar sensitivity to clutter and geometry, compressed engagement timelines, and the probabilistic nature of endgame lethality. Each event carried substantial cost—roughly $20 million per missile, $6.5 million in range operations, $1 million in instrumentation and telemetry, and $6–$11 million in ship operations, while also exposing ships, sailors, civilians, and collection assets to operational risk. Despite this investment, physical testing inherently under-sampled the performance space, leaving many edge cases and emergent behaviors unexplored.

Based on the technical realities we faced, the team adopted a structured alternative built around a reduced set of anchor tests supported by a high-fidelity modeling and simulation (M&S) environment. Although terms like "anchor tests" and "runs for the record" would later be formalized in DoD policy, the logic was already clear: use a small number of well-chosen physical events to calibrate and validate the digital environment, then rely on that environment to explore the broader performance space. We selected six anchor tests, replacing the original 12, to span radar performance boundaries, threat kinematic extremes, environmental stressors, and weapon-system timing constraints.

Once the models were verified, validated, and accepted by OPTEVFOR and Director, Operational Test and Evaluation (DOT&E), they were authorized for formal evidence in the Integrated Master Test Plan (IMTP) for Monte Carlo exploration of thousands of engagement variations, for excursions inside and outside the contracted envelope, and for regression testing supporting software builds and block upgrades. Although DoDI 5000.61 and the 5000-series T&E policies that explicitly permit M&S substitution would come later, the AEGIS BMD team was already demonstrating the feasibility and rigor of this approach.

The planning methodology we used closely resembled what would eventually become the Model-Based Test and Evaluation Master Plan (MBTEMP). Mission engineering defined operational threads; system models mapped those threads to system behaviors; test objectives were traced to model elements; anchor tests calibrated the model; and the validated environment enabled expanded exploration. Similarly, although the Integrated Test and Evaluation Management System (iTEMS) had not yet been conceived, our disciplined management of test artifacts, data pedigree, schedule dependencies, and configuration control mirrored the capabilities iTEMS now provides. In retrospect, the AEGIS BMD baseline effort stands as an early pathfinder that demonstrated the operational value of digital-engineering principles years before the DoD formally codified them.

As shown in Table 2, this approach produced measurable reductions in actual test costs. The accredited M&S environment enabled test-point reduction by identifying redundant conditions, applying design-of-experiments techniques to maximize coverage, and focusing physical tests on high-value regions, reducing the requirement from twelve planned tests to six anchor tests, a 50% reduction. Once validated, the models became reusable assets for software regression, block upgrades, threat updates, and mission-level analysis, extending their value far beyond the baseline delivery.

Table 2. AEGIS BMD Flight Test Cost Comparison: Original 12-Test Plan vs. Revised 6-Anchor-Test Plan

Cost Element	Per-Event Cost	Original Plan (12 Tests)	Revised Plan (6 Tests)	Savings
Missile / Test Article	~$20M	~$240M	~$120M	$120M
Range Operations	~$6.5M	~$78M	~$39M	$39M
Instrumentation & Telemetry	~$1M	~$12M	~$6M	$6M
Ship Operations & Crew	~$8 -11M (avg. $9.5M)	~$114M	~$57M	$57M
Safety / Risk Exposure	Non-monetary but significant	12 high-risk events	6 high-risk events	50% reduction in operational risk
Scenario Coverage	Limited	12 physical scenarios	Thousands via M&S	Massive increase
Total Quantifiable Cost	—	**~$444M**	**~$222M**	**~$222M saved (≈50%)**

Summary of ROI

- **Direct cost avoidance: ~$222 million**
- Operational risk reduction: 50% fewer high-risk live-fire events
- Scenario coverage: Orders of magnitude increase through Monte Carlo simulation
- Long-term asset creation: Accredited M&S environment reused for upgrades, regression, and threat updates
- Life cycle ROI: Compounding savings across future block upgrades

The AEGIS BMD baseline effort demonstrates that model-based T&E is not merely a cost-avoidance mechanism, it is a force multiplier for safety, operational confidence, and long-term modernization agility. Reducing twelve planned flight tests to six anchor events produced more than $222 million in direct savings, but the deeper return came from reducing high-risk live-fire evolutions by half, expanding scenario coverage from a dozen physical shots to thousands of analytically defensible engagement variations, and generating a validated M&S environment that continues to pay dividends across software builds, block upgrades, and threat-model updates.

Most importantly, this work, executed years before the DoD formalized MBTEMP, iTEMS, or digital-engineering policy, showed that the principles of DE deliver measurable, repeatable ROI even in the absence of a fully mature enterprise environment. When authoritative models, telemetry-anchored validation, and mission-thread traceability are combined, programs achieve higher confidence at lower cost, with greater safety and dramatically improved test sufficiency. This is the essence of digital-engineering value, and it remains as relevant today as it was when this pathfinding effort first proved what was possible.

Scenario 2—Applying DE Principles Forward: The Threat Digital Twin ATD

Life Cycle Cost Reduction Through DE and Threat Digital Twins

Scenario 2 presents the Threat Digital Twin ATD—a current program supported by Kitty Hawk Technologies—as a forward application of the principles demonstrated in Scenario 1. The AEGIS BMD baseline establishes the empirical proof of concept; the Digital Twin ATD illustrates how those same principles are being embedded in a current program from inception, and outlines the evidence framework that will quantify returns as the program matures. Program-sensitive cost data appropriately remains within the program office and is not presented here.

The AEGIS BMD baseline demonstrated what disciplined application of model-based principles can achieve in a mature ACAT I environment. The natural question that follows is whether those same principles translate forward into current programs, emerging threat domains, and mission areas where no adequate physical test or training asset previously existed. The Threat Digital Twin Advanced Technology Demonstration (ATD), a current effort supported by Kitty Hawk Technologies, provides a compelling answer.

Unlike AEGIS BMD, the Threat Digital Twin ATD is not yet at a stage where life cycle cost savings can be fully quantified. What can be described, and what makes this case analytically valuable, is the degree to which the program has deliberately embedded the foundational DE principles that produced measurable ROI in the AEGIS BMD case: authoritative digital baselines, physics-informed modeling, configuration-controlled artifacts, and a digital thread linking design, test, and validation. The argument is not that savings have been measured yet. The argument is that the structural conditions for measurable ROI have been deliberately established, and that the evidence framework for capturing those returns is already in place.

The Threat Digital Twin is not a generic visualization tool. It is a validated, physics-based, time-space correlated representation of a complex separating aerial vehicle, incorporating AI/ML-enhanced physics-informed neural network (PINN)-based threat motion modeling, six-degrees-of-freedom (6DoF) accounting for steady and unsteady aerodynamics, dynamic time and space-correlated RF signatures, and digital SE processes ensuring threat validity. These features make it a high-fidelity, configuration-controlled asset reusable across design, manufacturing, test, and sustainment, precisely the type of digital artifact the AEGIS BMD experience showed produces compounding life cycle returns.

Where the ROI Potential Lies

The AEGIS BMD case showed that the largest returns emerged from three mechanisms: reduction in physical test burden, creation of a reusable accredited M&S environment, and elimination of repeated re-characterization costs across software builds and block upgrades. The Threat Digital Twin ATD is structurally positioned to realize all three.

First, in a mission area where no adequate physical test asset previously existed, the threat digital twin directly substitutes for physical test articles across a range of design, analysis, and training scenarios. The cost avoidance associated with not fabricating multiple physical threat surrogates, each carrying substantial manufacturing, instrumentation, and range costs, is real, even if not yet fully tallied.

Second, the PINN-based dynamics model, once validated against flight-test data, becomes a reusable asset. New threat variants, geometry changes, and separation timing modifications can be assessed within the digital environment without new physical fabrication and extensive testing. This mirrors exactly the reuse value the AEGIS BMD accredited M&S environment delivered across subsequent block upgrades and regression test cycles.

Third, the authoritative digital baseline prevents the configuration drift that drives reverse engineering and re-characterization costs in legacy programs (DoD, 2025). Every manufactured unit, every design update, and every test configuration references the same validated source of truth, eliminating the ambiguity that compounds into life cycle technical debt. Once validated, the threat digital twin becomes a reusable asset that can be updated, extended, or adapted to represent new threat variants without requiring new physical test articles.

Evidence Framework for Quantifying ROI

A credible ROI case for the Threat Digital Twin ATD will be built from the following evidence streams, which the program is positioned to collect as it matures:

- Fabrication and test cost comparisons between physical threat surrogates and digital twin-enabled alternatives
- Reduction in re-characterization events as the validated model replaces repeated physical testing
- Configuration variance tracking across manufactured units against the digital baseline
- Upgrade cycle timelines for new threat variants developed within the digital environment versus those requiring new physical articles

These are not speculative metrics. They are the same categories of evidence that, in retrospect, would have quantified the AEGIS BMD ROI had the program been tracking them systematically from the outset. The AEGIS BMD scenario shows that that programs should instrument their DE adoption from the beginning, not reconstruct the evidence afterward.

Positioning in the Broader Argument

The Threat Digital Twin ATD case reinforces the central thesis of this paper in a specific and important way. The AEGIS BMD baseline proved that DE principles deliver measurable ROI even when adopted without a fully mature Digital Engineering Environment (DEE). The Threat Digital Twin ATD demonstrates that current programs, informed by that experience and supported by more sophisticated toolchains, are applying those principles earlier, more deliberately, and with a clearer eye toward life cycle returns.

The trajectory from AEGIS BMD to the Threat Digital Twin ATD is not merely historical. It is the institutionalization of hard-won lessons into current practice, which is precisely what the DoD's Digital Engineering Strategy calls for, and precisely what this paper argues the Department should systematically enable and resource. Importantly, none of these practices require full DEE maturity. The Threat Digital Twin ATD shows that even partial adoption, authoritative models, and a digital-threaded V&V workflow, establishes the structural conditions for measurable, repeatable ROI across the life cycle.

Cross-Cutting Analysis

A central insight emerging from both scenarios is that DE delivers measurable returns on investment across the entire acquisition life cycle, but the mechanisms, magnitude, and timing of those returns differ depending on when adoption occurs. Early adoption provides the greatest architectural freedom and the cleanest opportunity to establish a seamless digital thread from concept through sustainment. Yet the case examples presented here, particularly the AEGIS BMD baseline, demonstrate that mid-stream and late adoption often generate the most tangible, defensible, and operationally meaningful returns. This is because the costliest phases of the life cycle occur after Milestone C, when sustainment, obsolescence, regression testing, and configuration drift dominate program expenditures (DoD, 2025).

"It Is Never Too Late" to Adopt DE

One of the most persistent misconceptions in the DE community is the belief that DE must begin early, ideally at MSA or TMRR, to be effective. The empirical evidence from Scenario 1 directly contradicts this. The AEGIS BMD baseline adopted model-based principles at the end of EMD, at a point when the system was already transitioning into legacy sustainment. Despite this late entry, the program achieved a 50% reduction in physical flight tests, avoided hundreds of millions of dollars in cost, and created a reusable, accredited digital environment that supported future upgrades and regression testing. Perhaps most importantly, the digital environment expanded the quality and breadth of evidence available to decision-makers, far beyond what live testing alone could provide, by enabling thousands of analytically defensible engagement variations through Monte Carlo simulation.

This is not merely a case study; it is a proof point. Programs facing legacy technical debt, configuration drift, obsolescence pressures, and high sustainment costs stand to benefit the most from DE precisely because these are the challenges where traditional document-centric approaches fail. Scenario 2 reinforces this by demonstrating that current programs applying DE principles from inception are establishing the structural conditions for returns that will compound over the system's service life.

Governance and Metrics That Matter

Across both scenarios, the common denominator of successful DE adoption is not tool sophistication but governance discipline. DE succeeds when programs treat models and digital artifacts with the same rigor historically reserved for software—configuration control, formal verification and validation pipelines, and cross-functional governance structures that ensure digital artifacts remain authoritative and credible across life cycle phases.

Effective governance includes:

- Model configuration control with the same discipline applied to software baselines
- A formal VV&A pipeline for any model used to support decisions, test reduction, or certification
- A DE Steering Group integrating engineering, T&E, logistics, cybersecurity, and PMO functions
- Life cycle–aligned metrics, including:
    - Reduction in engineering change proposals (ECPs)
    - Reduction in physical test points
    - Decrease in sustainment downtime
    - Faster upgrade cycles
    - Avoided reverse engineering and re-baselining

Without governance, DE devolves into a collection of disconnected models. With governance, it becomes a strategic asset that compounds in value over time.

Risk Considerations

DE adoption is not without risk. Programs may over-promise benefits, underestimate the VV&A burden, or become dependent on proprietary toolchains that limit interoperability. Data pedigree issues can undermine model credibility, and workforce readiness remains a persistent challenge across the DoD. However, the AEGIS BMD example demonstrates that these risks are manageable even in the absence of modern DE infrastructure.

The most effective mitigation strategies include:

- **Incremental adoption**—focusing first on high-value models and authoritative baselines
- **Transparent assumptions**—rigorous documentation of model limitations
- **Telemetry-anchored calibration**—which builds credibility and reduces VV&A friction
- **Strict configuration control**—ensuring digital artifacts do not drift from physical reality
- **Cross-functional governance**—preventing DE from becoming siloed within engineering

Synthesis: What These Scenarios Tell Us About DE's True Value

The cross-cutting insight from both scenarios is that DE is not a monolithic capability but a portfolio of practices that deliver value in different ways at different times. Early adoption maximizes architectural flexibility and reduces integration risk. Mid-stream adoption improves verification, configuration control, and test sufficiency. Late adoption reduces sustainment burden, obsolescence risk, and regression test cost. In every case, the value proposition is tied to life cycle economics: when 60%–70% of total cost resides in O&S, even modest improvements in sustainment efficiency dwarf the cost of DE adoption.

The scenarios also demonstrate that DE's most powerful contributions extend beyond cost avoidance to improvements in safety, evidence quality, and operational confidence. Reducing live-fire events reduces risk to ships and sailors. Expanding scenario coverage improves decision quality. Creating authoritative digital baselines reduces rework and accelerates modernization. These benefits compound over time, creating a virtuous cycle in which each block upgrade becomes faster, cheaper, and more predictable than the last.

Conclusions and Recommendations

The analysis presented in this paper demonstrates that DE delivers measurable return on investment across both test and evaluation and life cycle management, and that these benefits do not depend on early adoption. Scenario 1 draws on the firsthand experience of John G. Fiore as Deputy Technical Director and Chief Engineer for AEGIS BMD, providing empirical evidence that model-anchored T&E reduced physical flight tests by 50% and generated cost avoidance of approximately $222 million, despite being implemented at the end of EMD and without a fully mature DE Environment. Scenario 2 draws on the technical work of Glenna Miller at Kitty Hawk Technologies, illustrating how the Threat Digital Twin ATD is applying those same principles from inception and establishing the structural conditions for measurable life cycle ROI through authoritative digital baselines, physics-informed modeling, and configuration-controlled digital threads. Together, these scenarios show that DE is not merely a design-phase enhancement but a life cycle capability that improves affordability, test sufficiency, and modernization agility regardless of when adoption begins.

The broader implication is difficult to ignore. Over the past decade, DoD policy has aggressively mandated DE adoption without first establishing a quantitative foundation for ROI or providing the resources required to implement it. Policy has outpaced evidence, and the burden of reconciling this gap has fallen disproportionately on programs least able to absorb it—particularly legacy programs where technical debt is highest and discretionary budgets are most constrained.

Correcting this imbalance requires a shift in how the Department conceptualizes and funds DE:

- **Treat DE as a program-level investment, not an unfunded mandate.** Resources must be aligned to the phases where the greatest returns are achievable.
- **Establish standardized ROI measurement frameworks** across both life cycle and T&E domains, enabling programs to quantify benefits in a defensible and repeatable manner.
- **Develop authoritative exemplars**—such as the AEGIS BMD baseline—that demonstrate how disciplined application of DE principles produces evidence-based value even without full DEE maturity.
- **Prioritize model credibility, configuration control, and data pedigree** as foundational elements of any DE effort, ensuring digital artifacts remain authoritative and reusable across decades of service life.
- **Support incremental adoption pathways** that allow programs to begin with high-value, low-risk digital practices rather than attempting enterprise-level maturity in a single leap.

Ultimately, DE is not a slogan, a compliance exercise, or a tool acquisition strategy. It is a disciplined engineering approach that reduces risk, accelerates modernization, and lowers life cycle cost when applied with rigor and supported with appropriate governance. The scenarios in this paper show what is possible when DE is implemented thoughtfully and pragmatically. The next step is institutionalizing these lessons so that every program, regardless of age, phase, or legacy burden, can realize the full life cycle value that DE promises.

Disclaimer: *The views expressed in Scenario 1 are solely those of John G. Fiore, based on his direct experience as the former Deputy Technical Director and Chief Engineer for AEGIS Ballistic Missile Defense. The technical content and analysis presented in Scenario 2 reflect the work and professional judgment of Glenna Miller in her capacity at Kitty Hawk Technologies. The views expressed in this paper do not represent the official position of the Department of Defense, the U.S. Navy, the Missile Defense Agency, or any other government agency or contractor.*

References

Attaran, M., & Celik, B. G. (2023). Digital twin: Benefits, use cases, challenges, and opportunities. *Decision Analytics Journal, 6*, 100165. doi:10.1016/j.dajour.2023.100165

DoD. (2018). *Department of Defense digital engineering strategy.* Office of the Deputy Assistant Secretary of Defense for Systems Engineering.

DoD. (2020). *Test and evaluation* (DoD Instruction 5000.89). Office of the Under Secretary of Defense for Research and Engineering.

DoD. (2022). *DoD modeling and simulation verification* (DoD Instruction 5000.61). Office of the Under Secretary of Defense for Research and Engineering.

DoD. (2023). *Digital engineering* (DoD Instruction 5000.97). Office of the Under Secretary of Defense for Research and Engineering.

DoD. (2025). *Operating and support cost-estimating guide.* Office of the Secretary of Defense.

GAO. (2023). *Leading practices: Iterative cycles enable rapid delivery of complex innovative products* (GAO-23-106222).

GAO. (2024a). *Weapon system sustainment: DOD identified operating and support cost growth but needs to improve the consistency and completeness of information to Congress* (GAO-24-107378).

GAO. (2025). *Weapon systems annual assessment: DOD leaders should ensure that newer programs are structured for speed and innovation* (GAO-25-107569).

GAO. (2025). *High-risk series: Critical actions needed to urgently address IT acquisition and management challenges (GAO-25-107852)*.

Whitehead, N. P., Light, T., Luna, A., & Mignano, J. (2024). *A framework for assessing the costs and benefits of digital engineering.* RAND.

PANEL 19. DATA-DRIVEN ACQUISITION MANAGEMENT: PERFORMANCE-BASED INCENTIVES, INDUSTRIAL CAPACITY CONSTRAINTS, AND WORKFORCE COMPETENCY ALIGNMENT

Thursday, May 7, 2026, 1315 – 1430 ET (1015 - 1130 PT)

Panel Summary:

True warfighting overmatch requires a fundamental realignment of the "Hexagon of Deterrence"—balancing the competing constraints of scope, schedule, cost, quality, resources, and risk to revitalize the U.S. defense industrial base. This panel explores how the Department can move beyond transactional compliance toward high-performance, Outcome-Based Contracting (OBC) that incentivizes supplier innovation and creative problem-solving. By integrating these results-driven partnerships with Portfolio Management competency standards that align with industry best practices, these research initiatives provide the framework for a more agile, transparent, and mission-aligned acquisition enterprise.

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)

Panel Presenters:

Outcome-Based Contracting: What Works, What Doesn't, and What's Next – *Daniel J. Finkenstadt, PhD, Commerce and Contract Management Institute*

The Value Hexagon of Warfighting Acquisition: Applying the Six Competing Constraints to Accelerate Warfighting Capability and Revitalize the U.S. Defense Industrial Base – *CAPT Jeffrey Dunlap, USN (Ret.), Naval Postgraduate School*

Portfolio Management Competency Standards – *Professor Bob Mortlock, Ph.D., Naval Postgraduate School*

Bob Marion, LTG USA, (ret.)—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.

Outcome-Based Contracting: What Works, What Doesn't, and What's Next

Daniel J. Finkenstadt, PhD—is Vice President of Research and Senior Fellow at the CCM Institute. He holds faculty appointments at George Mason University and the Johns Hopkins University Carey Business School. His research focuses on defense acquisition reform, perceived value in government procurement, and AI governance in public sector contracting. He is the author of Supply Chain Immunity (Springer Nature) and Bioinspired Strategic Design (Productivity Press) and has published in the Harvard Business Review, California Management Review Insights, and numerous peer-reviewed journals. He is a graduate of Naval Postgraduate School (MBA 2011) and a previous assistant professor of defense management at NPS (2020–2023).

Tim Cummins—serves as the Executive Director of the CCM Institute and as the President of World Commerce & Contracting. Until 2023, he was a Professor in the Law School at the University of Leeds (UK). For 25 years, Tim has led research and the development of standards for the Commercial and Contract Management discipline, an achievement that was recognized by the Financial Times "Market Shaper of the Year" award in 2019 and again in 2025. Over that time, Tim led development of the world's only international not-for-profit association dedicated to commercial and contract management, working with almost 100,000 members to improve the quality and integrity of trading relationships in both public and private sector.

Abstract

This research addresses a foundational question for defense acquisition leaders, namely, what are outcome-based contracts (OBCs), under what conditions should the Department of War employ them, and what institutional capacities must be in place for them to succeed? Drawing on a multi-phase, mixed-methods research design that included a comprehensive literature review, 14 semi-structured interviews with senior commercial and contract management professionals across eight countries, two practitioner focus groups (N = 34), a federal acquisition community survey, and an executive roundtable with 62 senior acquisition leaders, this study integrates both U.S. federal and global commercial perspectives to identify five critical success factors for OBC implementation: requirements definition, data sufficiency, inter-party trust, governance capability, and oversight balance. The theoretical foundation integrates Graeber's (2001) anthropological theory of value, Zeithaml's (1988) perceived value framework, Vargo and Lusch's (2004, 2008) Service-Dominant Logic, and empirical research on perceived service quality in business-to-government settings (Finkenstadt, 2020). A central finding is that outcome-based strategy and outcome-based contracts are distinct constructs; conflating them produces implementation failure. The study offers five policy recommendations directed at defense acquisition leadership, including FAR repositioning, governance training investment, portfolio prioritization, and structured low-risk piloting mechanisms.

Keywords: outcome-based contracting, defense acquisition, performance-based contracting, value co-creation, contract governance, Federal Acquisition Regulation

Introduction

Interest in outcome-based contracting approaches is increasing across both defense and civilian agencies, yet confusion persists about what outcome-based contracts (OBCs) actually mean and how to implement them effectively. Agencies want measurable results, but the acquisition workforce lacks clear definitions, practical guidance, and the institutional infrastructure to support a shift from paying for inputs and activities to paying for demonstrable mission outcomes.

OBCs represent a fundamental shift from traditional procurement by prioritizing measurable results over rigid process requirements. This framework empowers suppliers to leverage innovative methods, emerging technologies, and creative delivery strategies to achieve defined outcomes. Such flexibility proves especially valuable in dynamic defense environments where requirements evolve rapidly and prescriptive specifications can hinder effective responses from both program offices and their industry partners.

The CCM Institute, the joint research arm of NCMA and WorldCC, led this research initiative with a clear objective: to clarify what OBCs are (and what they are not), to distinguish outcome-based strategy from outcome-based contract structure, and to identify the practical conditions that the Department of War must establish for successful OBC implementation. The research drew on the CCM Institute’s expertise in commercial contracting, governance frameworks, and global benchmarking, integrating perspectives from both U.S. federal acquisition professionals and senior commercial practitioners across eight countries to ensure that findings reflect the full range of institutional contexts in which outcome-based approaches are being adopted.

A key finding is that outcome-based strategy and outcome-based contracts are not the same thing. Outcome-based strategy is a broader approach to mission delivery that focuses on measurable impact. Outcome-based contracting is a specific contractual structure that links compensation to defined results. Confusing these two constructs leads to poor implementation. Agencies must first adopt an outcomes mindset before they attempt to embed outcomes into contract structures.

The U.S. Government has recently issued guidance in the FAR Companion Guide (Parts 11 and 37) that, while well-intentioned, risks limiting the OBC concept and confusing the workforce. This research offers a corrective path. The federal government has signaled strong interest in performance-driven procurement; however, without conceptual clarity and workforce readiness, outcome-based approaches risk being reduced to compliance language rather than meaningful reform.

This paper proceeds as follows. The literature review establishes the theoretical and empirical foundations for understanding value in contractual relationships. The methodology section describes the five-phase, mixed-methods data collection effort conducted from Spring 2025 through February 2026, which incorporated both U.S. federal acquisition and global commercial contracting perspectives. The findings section presents the five critical success factors that emerged consistently across all data sources. Two special topics, AI’s impact on OBC management and the urgent need for governance training, receive dedicated treatment. The paper concludes with five policy recommendations directed at defense acquisition leadership, a discussion of limitations, and a Phase II research agenda.

Literature Review

This section presents the theoretical and empirical foundations that guided the research design and informed the analytical framework. The review integrates anthropological theory of value, consumer perceived value research, Service-Dominant Logic, and empirical findings on perceived service quality in business-to-government procurement to establish a coherent lens for understanding why contracts can remain technically compliant while failing to deliver intended outcomes.

Contracts as Value Architecture

The anthropological theory of value developed by Graeber (2001) offers three propositions with direct implications for contracting practice. First, Graeber argued that value measures the importance of actions, not the properties of objects. When parties assign value to something, they are registering a collective judgment about which human actions matter. Applied to contracting, this means that every requirement, deliverable, incentive, and remedy reflects a judgment about relative importance.

Second, objects and documents function as condensed action-claims. They do not hold value intrinsically; they hold it because they encode, stabilize, and make portable a shared understanding of which actions were significant and why. A contract is therefore not the value itself; it is the vessel that preserves a shared understanding of value so that coordinated action can occur (Graeber, 2001).

Third, exchange is fundamentally communicative. When parties exchange goods, services, or obligations, they are sending a signal about what matters in their relationship. The terms of exchange are never neutral; they always express something about what both parties believe is worth doing. Every contract communicates a value hierarchy to everyone performing under it (Graeber, 2001).

These three propositions produce a single insight highly relevant to OBC practice. A contract is a structured communication about collective value. That communication either faithfully represents the underlying purpose of the agreement or it gradually replaces it with something else. This theoretical lens helps explain the persistent challenge identified across our research, the tendency for contracts to drift from purpose toward compliance, even when both parties intend otherwise.

Perceived Value in Procurement

Zeithaml's (1988) seminal research on consumer perceived value established that perceived value is the buyer's overall assessment of the utility of a product or service based on perceptions of what is received relative to what is given. Her research revealed four distinct conceptions of value operating simultaneously among buyers: value as low price, value as getting what the buyer wants, value as quality relative to price paid, and value as the total balance of benefits received against all sacrifices made.

The fourth and most integrative conception, encompassing all costs (monetary, time, effort, risk) against all benefits (functional, experiential, relational), aligns directly with Graeber's action-based view. For OBC practice, this framework reveals that traditional procurement's emphasis on Lowest Price Technically Acceptable (LPTA) represents only the first and narrowest conception of value. Outcomes-based approaches, by contrast, operate in the fourth conception, evaluating the full relationship between what is invested and what results are achieved (Zeithaml, 1988).

Service-Dominant Logic and Value Co-Creation

Service-Dominant Logic (SDL), developed by Vargo and Lusch (2004, 2008), fundamentally reframes value creation in ways directly relevant to government services and OBC design. Four foundational premises are particularly consequential. First, service is the fundamental basis of exchange; all economies are service economies, and the underlying exchange is the application of knowledge and skills for the benefit of another party. Second, value is co-created by multiple actors, always including the beneficiary; providers cannot deliver value unilaterally but can only offer value propositions, with value emerging through joint interaction as providers and beneficiaries integrate resources.

Third, value is always uniquely and phenomenologically determined by the beneficiary, meaning it is idiosyncratic, experiential, and contextual. Fourth, operant resources (knowledge and skills rather than physical resources or labor hours) are the fundamental source of strategic benefit (Vargo & Lusch, 2004, 2008).

SDL transforms the procurement vocabulary in consequential ways. Traditional procurement treats contractors as suppliers from whom the government extracts deliverables. SDL reframes the relationship as joint value creation. It would offer that the government and contractor together produce outcomes neither could achieve alone. This shift from "market to" to "market with" is particularly significant for OBCs, where the collaborative, adaptive partnership described in the FAR Companion Guide is not merely a contracting preference but a structural requirement for value realization (Vargo & Lusch, 2004, 2008).

Empirical Evidence on Quality, Value, and Procurement Choice

These theoretical frameworks receive direct empirical support from research on business-to-government (B2G) knowledge-based services. Research with more than 630 Department of War procurement professionals establishes a clear causal path from perceived quality through perceived value to procurement choice (Finkenstadt, 2020; Finkenstadt & Zeithaml, 2020).

Perceived service quality is a second-order construct comprising four dimensions: employee capability, intelligent solutions, dependability, and understanding of customer requirements. Perceived quality strongly predicts perceived value, explaining the vast majority of variance in value perceptions. Critically, perceived service quality attributes are more than twice as important as price in B2G procurement choice. The ability to provide intelligent solutions is the single most important attribute, more than twice as important as price (Finkenstadt, 2020; Finkenstadt & Zeithaml, 2020).

Choice-based conjoint analysis reveals substantial willingness to pay for quality improvements. When comparing offers, procurement agents would trade off up to a 41% price premium for high-confidence over low-confidence quality ratings. Agents would nearly always opt out entirely rather than select a low-confidence offer, even at the lowest price. This directly contradicts LPTA assumptions and has profound implications for how OBCs should be valued, structured, and evaluated (Finkenstadt, 2020).

Contracts as Three-Function Value Systems

Viewing contracts through this theoretical lens reveals that they perform three functions beyond their legal enforceability, each of which must work for the contract to transmit value effectively (Finkenstadt, 2026). First, contracts capture value in the form of requirements, deliverables, incentives, and remedies that define what counts as worthwhile performance. Second, contracts communicate value, providing the structure sends a signal to every party performing under it about what matters most. Detailed specifications communicate a preference for predictability and control; flexible performance metrics communicate trust and an orientation toward results. Third, contracts transfer value over time through payment schedules, acceptance criteria, data rights provisions, and risk allocation mechanisms that govern how value moves between parties as the agreement is executed.

When the contractual vessel leaks, when clauses accumulate that are disconnected from purpose or when incentives drift from the behaviors that produce real outcomes, the value escapes before it can be realized. This dynamic applies fully to physical deliverables, not just services or relational agreements. Technical specifications are not value-neutral: tighter tolerances prioritize reliability over cost; accelerated delivery prioritizes readiness over margin; redundancy requirements prioritize resilience over efficiency. Requirements are prioritized beliefs written in technical language (Finkenstadt, 2026).

The Contracting Spectrum and Narrative Drift

The contracting literature often presents a clean binary in which traditional contracts specify activity, outcomes-based contracts specify results. The reality is a spectrum, and where a contract sits on that spectrum determines how faithfully it transmits value (Finkenstadt, 2026). Traditional models specify tasks, inputs, and methods. The dominant signal to everyone performing under them is to follow the instructions, and value is located in compliance with process. Outcomes-based models specify the final effects that matter and tie payment and evaluation to whether those effects are achieved.

Performance-based contracting was designed to occupy the productive middle, and in principle it does. But in practice, performance-based models have proven vulnerable to a failure mode structurally similar to the one they were designed to correct, in which the measures become the mission. Providers optimize for the metric, not the effect the metric was originally meant to capture. This is not cynicism or gaming; it is the natural behavioral response to what the contract signals is important (Finkenstadt, 2026). The practical implication is that performance-based contracting is a necessary but not sufficient reform.

Federal Acquisition Regulation Definition

The recently released FAR Companion Guide defines outcome-based contracting as: *"Outcome-based contracting is a variation of performance-based contracting that emphasizes delivery of specific, defined outcomes through a collaborative, adaptive performance framework, rather than transactional delivery of specified services or products."*

The Guide further states: *"The essence . . . is transforming the government-contractor relationship from a transactional exchange to a strategic partnership unified around delivery of defined performance outcomes."*

Empirical Evidence on OBC Type Selection

The research findings support the importance of distinguishing between availability outcome-based contracts (aOBCs) and economic outcome-based contracts (eOBCs). This distinction has been empirically validated in a multi-industry survey of 259 buyers and sellers using OBCs in complex industrial services, which found that the two forms differ meaningfully in their benefit and risk profiles depending on contextual conditions (Bohm et al., 2016).

Availability OBCs are structured around ensuring the availability of a system, asset, or capability, with the contractor held accountable for readiness and uptime rather than discrete deliverables. Economic OBCs tie contractual outcomes to economic results such as cost savings, efficiency gains, or value delivered relative to investment. Both buyers and sellers attach significantly higher perceived benefits to eOBCs compared to aOBCs; however, on average, both forms are perceived as equally risky and perform equally well from both perspectives (Bohm et al., 2016).

The critical insight is that contextual conditions determine which form is superior. In technologically turbulent environments, buyers perceive significantly more benefits from eOBCs, making them the preferred option, and the corresponding risk increase for sellers in these conditions was found to be non-significant. When product innovativeness is high, aOBCs emerge as the better option because sellers associate significantly higher risks with eOBCs when they lack accumulated experience deploying innovative offerings in customers' specific processes (Bohm et al., 2016).

Illustrative Case: The F-35 Program

The F-35 Joint Strike Fighter program provides one of the clearest available illustration of what happens when contractual structure drifts from value. The program is the largest defense acquisition in history, with a projected lifecycle cost now estimated at over $2 trillion (GAO, 2025). The F-35's mission is air superiority, and its value to the warfighter is combat capability. For years, however, the program's contract incentive structure rewarded something different, specifically on-time delivery of aircraft.

According to a September 2025 Government Accountability Office (GAO) report, the program paid contractors hundreds of millions of dollars in incentive fees, even as contractors delivered engines and aircraft consistently late, and as the aircraft being delivered were frequently non-combat-capable. The structure allowed delivery up to 60 days late while still earning partial on-time delivery fees. By July 2024, the program had begun provisionally accepting 174 aircraft that lacked combat-capable software. As of early 2025, the F-35A fleet's mission-capable rate stood at approximately 52%, against an 80% target (GAO, 2025).

In the theoretical terms established above, the contract had stopped communicating the right value hierarchy. Delivery timing became the operative signal. Combat readiness, the actual purpose of the program, had been subordinated to production throughput. The GAO's analysis was direct: unless the program reevaluates its use of incentive fees and better aligns them with desired production outcomes, it will continue to reward contractors for underperformance (GAO, 2025). That is not a criticism of the contractors; it is a diagnosis of a contract that lost its value narrative.

Historical and International Precedents

Several historical and international cases reinforce the empirical foundations. Guajardo et al. (2012) analyzed 305 Rolls-Royce aircraft engines under the "Power by the Hour" model and demonstrated that OBCs were associated with a 25% to 40% increase in product reliability compared to traditional input-based contracting forms. Australian defense navy sustainment contracts focused on vessel readiness created a shared operational objective that strengthened collaboration and reduced compliance-oriented behavior. The United Kingdom's Social Value Act further embedded OBC principles through legislation integrating measurable social outcomes into public procurement. These cases confirm that outcome-based incentive alignment produces measurable performance improvements when the institutional conditions are in place.

Methodology

This study employed a multi-phase, mixed-methods research design to build a comprehensive evidence base from both the scholarly literature and the practitioner community, integrating U.S. federal acquisition perspectives with global commercial contracting experience. Data collection proceeded across five sequential phases from Spring 2025 through February 2026. The phased approach enabled iterative refinement of the research questions, with findings from each phase informing the design and focus of subsequent phases.

Phase 1: Literature Review (Spring 2025)

The research began with a comprehensive review of academic literature, practitioner publications, and policy documents on OBCs, performance-based contracting, and related frameworks. The review examined historical precedents and international examples to understand how outcome-based approaches have evolved in practice, and it helped clarify common incorrect assumptions and recurring implementation challenges.

The literature review also incorporated foundational theoretical work on the nature of value in contractual relationships. Drawing on Graeber's (2001) anthropological theory of value, Zeithaml's (1988) research on consumer perceived value, and the Service-Dominant Logic framework developed by Vargo and Lusch (2004, 2008), we examined the proposition that contracts function not merely as legal instruments but as structured communications about collective value, where every requirement, incentive, payment term, and performance metric constitutes a statement about what the parties believe matters. This theoretical grounding was further informed by empirical research on perceived service quality and value in business-to-government knowledge-based services (Finkenstadt, 2020; Finkenstadt & Zeithaml, 2020), which provided quantitative evidence on how procurement professionals actually weigh quality against price in selection decisions.

Together, these frameworks offered an analytical lens for understanding why contracts can remain technically compliant while failing to deliver intended outcomes, a pattern directly relevant to the challenges facing OBC implementation in the defense acquisition context. Insights from early focus group discussions and global expert interviews further shaped our understanding, highlighting concerns around unclear definitions, measurement risk, workforce readiness, and cultural variation in contracting practice. This allowed us to establish the core analytical categories (Requirements, Data, Trust, Governance, and Oversight) that guided the subsequent survey and roundtable. Our literature review was supplemented in later stages by the publication of the FAR Companion Guides as part of the Revolutionary FAR Overhaul that included a discussion of OBCs in both parts 11 and 37.

Phase 2: Global Expert Interviews (June–September 2025)

Between June 10 and September 3, 2025, the CCM Institute conducted 14 semi-structured interviews with senior commercial and contract management professionals to capture global perspectives on OBC practice. This phase was designed to complement the U.S.-focused focus groups and survey by gathering insights from practitioners operating under diverse legal, regulatory, and cultural contracting regimes. Interviewees were drawn from 14 contract management executives in eight countries: the United States, United Kingdom, Australia, Japan, Ireland, Canada, France, and Colombia. All held senior or executive-level positions, including Chief Legal Officer, Executive Vice President, Managing Director, Commercial Director, Head of Contract and Business Management, and senior leaders in procurement, supply chain, and contract management.

The interviews explored core topics including how practitioners define and distinguish OBCs from performance-based contracts, real-world implementation experiences across sectors (infrastructure, aerospace and defense, technology, consulting, and public services), perceived benefits and barriers, risk allocation and governance challenges, and the role of trust and organizational culture in sustaining outcome-oriented relationships. The geographic and functional breadth of this sample enabled us to identify patterns that transcend any single national procurement framework and to test whether the critical success factors emerging from U.S.-focused data held across diverse institutional contexts.

Phase 3: Focus Groups at NCMA World Congress (July 2025)

We organized two guided focus group sessions at NCMA World Congress in Grapevine, Texas, with 19 participants in the first group and 15 in the second, all senior U.S. procurement leaders. These sessions explored core topics including OBC definitions, hands-on implementation experiences, typical challenges, and success factors. The group setting encouraged participants to share contrasting viewpoints, question established thinking, and exchange practical lessons. This collaborative approach revealed important themes and conflicts that individual interviews might miss, offering essential practitioner perspectives that enriched the broader data gathering effort.

Phase 4: OBC Survey (Fall 2025)

On September 1, 2025, we deployed a comprehensive OBC survey within the federal acquisition community. This quantitative and qualitative research instrument aimed to confirm and build upon findings from the literature review, global expert interviews, and focus groups. The survey collected information about current OBC usage, perceived benefits, implementation obstacles, and areas requiring additional support. This approach enabled us to validate initial findings across a larger population and gain deeper insights into how practitioners actually experience outcome-based contracting methods in their daily work.

Phase 5: Executive Roundtable at NCMA Nexus (February 2026)

In February 2026, an Executive Roundtable was organized during NCMA Nexus in Atlanta, bringing together 62 senior acquisition executives to conclude the research phase. We shared initial findings to confirm major themes and validate preliminary insights. Discussions centered on actionable recommendations, implementation challenges, and organizational constraints. This forum provided a crucial opportunity to gather executive viewpoints, enhance the analysis, and ensure the report balanced strategic objectives with real-world operational considerations. This was the final opportunity to capture executive-level insights before report drafting.

Findings

This section synthesizes all data collected from Spring 2025 through February 2026 into a coherent narrative organized around five critical success factors and two special topics. These factors consistently emerged across the literature, global expert interviews, focus groups, survey, and roundtable discussions, forming a consolidated, evidence-based foundation for effective OBC design and execution that reflects both U.S. federal acquisition and international commercial contracting perspectives. The core message is clear. If the answer is “no” on any of these five factors, the OBC model will struggle.

Factor 1: Requirements

The first and most fundamental question is whether the desired outcomes can actually be articulated in a way that is measurable, attributable, and contractually actionable. Our research documented a significant evolution over the eight-month data collection period. Early participants struggled to distinguish OBCs from PBCs, with confusion around what constitutes an “outcome” versus a “performance metric.” By the roundtable phase, the discourse had shifted from definitional confusion to institutional design complexity.

Practitioners now focus on more sophisticated questions. They ask what parts of outcomes are genuinely within supplier control, what buyer actions materially affect success, and how uncertainty should be structured within the requirements framework. The theoretical framework reinforces this finding. When requirements fail to capture the true value hierarchy, when they specify proxy activities rather than the outcomes those activities are meant to produce, the contract communicates the wrong priorities to everyone performing under it. As the F-35 case illustrates, even well-intentioned requirements can encode priorities that diverge from the program’s actual mission (GAO, 2025; Graeber, 2001).

The global expert interviews reinforced this finding while adding important nuance around the spectrum of requirements precision. Interviewees distinguished between two OBC types: availability OBCs, which feature clearly defined and measurable outcomes such as 95% system uptime, and economic OBCs, which involve broader goals that are harder to quantify, such as reducing homelessness by 20%. This distinction maps directly to the requirements challenge. International practitioners noted that ambiguity in defining success metrics was the most common source of contractual disputes, and that external factors often influence outcomes beyond supplier control, making requirements definition inseparable from the governance and trust factors discussed below.

Factor 2: Data Sufficiency

Measurement difficulty remains a persistent barrier, but our research documents greater structural clarity around performance management. Practitioners now articulate specific design choices including clear outcome definition (both quantitative and qualitative), success thresholds (minimum, maximum, and stretch), measures of effectiveness defined in the RFP, protocols for handling unforeseen conditions, milestone approaches, RACI models, and two-way performance accountability mechanisms.

The maturity has shifted from “measurement is hard” to “measurement requires structured design choices and governance discipline.” The empirical research on B2G procurement reinforces the stakes. When meaningful quality signals exist, procurement professionals demonstrate a strong capacity to distinguish value (Finkenstadt, 2020; Finkenstadt & Zeithaml, 2020). The data infrastructure for OBCs is not merely an administrative requirement; it is the mechanism through which value becomes visible and consequential.

Global interviewees corroborated this progression while highlighting a persistent structural barrier: baseline data is frequently unavailable or unreliable, particularly in sectors or regions where outcome-based approaches are being attempted for the first time. Several interviewees noted that without credible baselines, outcome targets become arbitrary, undermining the legitimacy of the entire measurement framework. International practitioners also identified AI and data analytics as increasingly important enablers of outcome measurement, with emerging applications in automated reporting against outcome metrics, predictive analytics for performance trends, and enhanced data analysis for measurement frameworks. The Schiphol Airport lighting-as-a-service contract with Philips was cited as an exemplar where continuous performance monitoring made data sufficiency a built-in feature of the commercial model rather than an afterthought. However, interviewees cautioned that technology alone does not resolve the underlying design challenge noting that organizations must first define what constitutes a meaningful outcome before instrumenting its measurement.

Factor 3: Inter-Party Trust

Risk misalignment remains central to OBC challenges, with recognition that outcomes are often influenced by factors beyond supplier control. Practitioners now draw clear distinctions between products (where the supplier exercises higher control) and services (where control is shared), while recognizing that supplier control depends significantly on buyer readiness and governance capability.

Empirical research on OBC risk perception validates this nuanced view. Bohm et al. (2016) found that OBCs generally tend to shift risk toward the seller; however, on average, sellers do not perceive economic OBCs as more risky than availability OBCs, suggesting that the increased operational responsibility is offset by the stronger incentive alignment and the greater flexibility eOBCs provide. The critical exception occurs when product innovativeness is high, where sellers associate significantly higher risks with eOBCs because they lack accumulated experience deploying innovative offerings in customers’ specific processes. For defense acquisition, this means that the trust factor must be assessed relative to the maturity and complexity of the underlying systems, not assumed from the contract form alone, pointing again to outcomes-based requirements as a first principal for success.

Empirical evidence from B2G research further illuminates this dynamic. In knowledge-based services, collaborative sharing behaviors, including communication, feedback, and teaming, are highly correlated with perceptions of service quality, while unilateral customer-contribution behaviors are not (Finkenstadt, 2020). This suggests that in OBC contexts, the collaborative dimension of value co-creation is more closely tied to quality perceptions than the extractive dimension. Firms and agencies that invest in the working relationship are more likely to realize the value OBCs are designed to produce (Vargo & Lusch, 2008; Finkenstadt, 2020).

The global expert interviews added a critical dimension to the trust factor by surfacing the financial and cultural barriers that erode trust before performance even begins. Interviewees consistently identified delayed payments and cash flow challenges as a structural impediment to supplier willingness, particularly for small and medium providers who lack the capital reserves to absorb the financial risk inherent in outcome-contingent payment models. As one interviewee candidly observed, "Human beings are reckless, and human beings run companies," highlighting how even well-designed contractual structures break down in practice when human judgment and organizational incentives are misaligned. Risk pricing emerged as a persistent dilemma: interviewees reported that outcome-based risk allocation often leads to either inflated costs, as suppliers price in uncertainty, or dangerous underbidding, as competitors gamble on favorable conditions. The international sample also revealed meaningful cultural variation in how trust operates within OBC relationships. Western contracting traditions tend to front-load trust through detailed legal specification, while practitioners from Japan and Colombia described relational contracting norms that embed trust in long-term partnership expectations and iterative collaboration. These cultural differences do not invalidate the trust factor but suggest that its operationalization must be calibrated to the institutional and relational context of each engagement.

Factor 4: Governance Capability

One of the clearest evolutions documented by our research was the growing emphasis on governance as distinct from traditional contractual clauses. Participants acknowledged that acquisition professionals default to rigid terms and conditions, that escalation structures and adaptive governance are poorly understood, and that OBCs require collaborative decision-making frameworks rather than compliance-heavy control mechanisms.

The theoretical framework helps explain this dynamic: governance is the institutionalized process for maintaining shared meaning, for allowing the social process that produces value to continue throughout performance rather than calcifying at award. Joint governance boards, periodic performance reassessment, data transparency mechanisms, and structured renegotiation triggers are not administrative overhead; they are essential to preserving the contract's value narrative. SDL literature reinforces that, because value is co-created and phenomenologically determined by the beneficiary, governance mechanisms must allow both parties to continuously recalibrate what "success" means as circumstances evolve (Vargo & Lusch, 2008).

The CCM Institute's relational contracting and governance research provides a practical architecture for operationalizing these principles, emphasizing joint working structures, defined meeting cadence, communications protocols, escalation paths, decision rights, and continuity mechanisms when key personnel change (NCMA, 2025). Critically, this framework stresses that governance readiness must be assessed on both sides of the relationship. Outcome-based models fail when the buyer does not fulfill its own obligations—including timely data access, responsive decision-making, consistent payments, operational cooperation, and stable governance resourcing. A readiness gate that tests not only "can the supplier deliver" but also "can the agency govern" should be a precondition for any OBC commitment.

Global interviewees identified governance capability as the factor most consistently underinvested across all regions and sectors examined. Several noted significant loss of institutional knowledge in the post-COVID period, as experienced contract managers retired or changed roles during the disruption, leaving organizations without the tacit expertise needed to manage collaborative relationships. Interviewees also emphasized that maintaining a collaborative mindset during contract challenges, when disputes arise over outcome attribution or when external conditions shift the performance baseline, is the hardest governance discipline to sustain. The international perspective revealed that organizations with mature governance practices, particularly those in the UK and Australia with experience in availability-based defense sustainment models, have developed structured change management processes that explicitly anticipate the need to renegotiate outcome definitions during performance. These practitioners described governance not as oversight but as the active stewardship of a shared commercial relationship, a framing that aligns closely with the SDL concept of value co-creation.

Factor 5: Oversight Balance

The roundtable revealed a fundamental tension between commercial flexibility (post-award refinement) and the government's preference for full pre-award definition. Current findings emphasize the need for adaptive governance and joint accountability while managing the risk that wrong metrics can distort behavior, especially in complex defense environments.

Artificial Intelligence (AI) integration introduces new complexity, but participants demonstrated practical realism about current limitations, emphasizing the need for human-in-the-loop decision authority and workforce capability across builders, users, and validators. The oversight factor embodies the broader challenge of maintaining a contract's original value narrative over time, a challenge that requires leaders to ask different diagnostic questions, ones that test whether the agreement will actually transmit the value it was designed to carry.

The global interviews surfaced additional oversight challenges that extend beyond the U.S. regulatory context. Interviewees noted that budget structures in many organizations, both public and private, cannot easily accommodate the variable payment schedules that pure OBCs require. Legal frameworks designed for traditional contracting often impose rigidity at precisely the points where OBCs demand flexibility, creating a structural tension between accountability and adaptability. Perhaps the most significant finding from the international sample was the near-universal preference for hybrid models: pure OBCs remain rare in practice, with most organizations incorporating outcome-linked elements within otherwise traditional contract structures. As one interviewee summarized, "If I can generate $100 of value, I get $20, and the client keeps $80—that's aligned." This gain-share logic, in which baseline costs are covered through traditional mechanisms while incremental value is shared through outcome-contingent payments, was described as the most practical and sustainable approach to outcome-based oversight across every region and sector represented in the study.

Discussion: AI's Impact on OBC Management

AI is emerging as a transformative force in OBC management, though adoption remains cautious and strategic. Current applications focus on enhanced data analysis for outcome definition, accelerated solution development, and improved market research capabilities. AI tools show particular promise in automated reporting against outcome metrics, predictive analytics for performance trends, natural language processing for contract analysis, and enhanced data analysis for measurement frameworks.

These applications align with broader trends. The CCM Institute's AI in Contracting 2026 survey of 518 global contracting professionals found that practitioners see the greatest value for AI in risk assessment and compliance (65%), contract performance monitoring (60%), and contract generation (55%). Interest drops sharply for negotiation support (29%), indicating that judgment, context, and relationship management remain firmly in the domain of human expertise.

A critical challenge for AI-enabled OBC management is not the sophistication of the AI tools themselves but the quality and intentionality of the data they rely on. The GDSD (Goals, Decisions, Signals, Data) model developed by Finkenstadt et al. (2022) argues that in a data-saturated environment, organizations must resist the instinct to begin with observation and instead start with the intended outcome. The model's logic proceeds in four steps: define the goal, identify the decisions required to achieve it, determine what signals would inform those decisions, and only then specify the data needed to generate those signals. Data that does not connect to a decision-relevant signal is noise, regardless of how easy it is to collect.

Practitioners emphasize critical limitations. Security and privacy remain the top barrier to AI adoption, cited by 68% of respondents. Data output quality, including the risk of AI-generated errors and hallucinations, is the second-highest barrier at 55%. Over-reliance on AI emerged as the most pronounced area of worsening concern, reinforcing fears that human judgment is being diluted rather than augmented. For OBC management, where outcome interpretation inherently involves subjective assessment and collaborative negotiation between buyer and supplier, this risk is particularly acute.

The risk of over-reliance is compounded by what Tillipman (2026) identifies as automation bias in the federal procurement context. She defines this bias as the tendency for acquisition professionals to defer to AI-generated outputs rather than independently evaluate them, particularly under severe time and resource pressures. A related risk involves what Tillipman (2026) terms "nested opacity," whereby the AI tool's architecture, training data, and internal decision logic are frequently shielded by commercial licensing terms. For OBCs, if an agency relies on a vendor's proprietary AI system to assess whether contract outcomes have been achieved, and the basis for that assessment cannot be reconstructed or challenged, the collaborative performance review process that distinguishes OBCs from compliance-driven models is undermined at its foundation. Agencies deploying AI for OBC performance monitoring should secure contractual rights sufficient to understand how the tool applies outcome criteria and to question, override, or reject its outputs when governance judgment warrants (Tillipman, 2026).

Discussion: The Urgent Need for Contract Governance Training

The transition to outcome-based contracting exposes a critical capability gap in the defense acquisition workforce, specifically the fundamental distinction between contract management and contract governance. Contract management focuses on administrative compliance and transactional oversight. Contract governance demands strategic relationship management, adaptive problem-solving, and collaborative outcome interpretation throughout the contracting life cycle. This distinction is poorly understood across government contracting professionals, creating a workforce development crisis that threatens OBC success.

Our research reveals that administrative performance measures dominate. Cycle time emerges as the top performance measure at 50%, while monitoring customer satisfaction ranks at 45% and supplier satisfaction at 40%. However, critical governance capabilities receive significantly less attention. Collaborative compliance monitoring stands at only 20%, and joint risk management at just 15%, highlighting the administrative bias in current practices.

OBCs require contracting professionals to master entirely new competencies such as managing contractual ambiguity, facilitating joint problem-solving sessions, interpreting outcome data collaboratively with vendors, and maintaining strategic alignment as requirements evolve. The theoretical insight that governance mechanisms are not administrative overhead but rather institutionalized processes for maintaining shared meaning throughout contract performance reinforces the urgency. Without governance capability, the value narrative that animates an OBC at award will inevitably drift toward compliance-driven proxies. Addressing this gap requires more than generic training; it demands role-based capability mapping that identifies distinct competency needs for contracting officers, program managers, requirements owners, financial oversight staff, and industry partners. CCM Institute’s Contract Management Standard 4 (CMS4) offers a structured competency framework that can serve as a starting point for this effort, providing defined capability levels across the commercial functions that OBC governance demands.

A Derived Definition of Outcome-Based Contracts

Based on our research, we propose the following definition. *“The essence of outcome-based contracts lies in shifting focus from ‘how work is done’ to ‘what results are achieved.’ Suppliers are given flexibility in methodology and approach while being held accountable for delivering agreed-upon outcomes through clear performance metrics, risk-sharing arrangements, and adaptive frameworks that encourage innovation and continuous improvement. Payment structures are contingent upon successful delivery of these outcomes, creating aligned incentives that drive both parties toward shared strategic objectives.”*

This definition demonstrates strong conceptual alignment with the updated FAR Companion Guides in their shared emphasis on results-oriented contracting. However, there are important distinctions in scope, implementation mechanisms, and regulatory context. Both the proposed definition and FAR guidance emphasize shifting from prescriptive instructions to outcome-focused requirements, supplier flexibility and innovation, and performance accountability. Key contrasts include the proposed definition's explicit embrace of contingent payment structures (which the FAR guidance positions within existing federal payment frameworks), more collaborative risk distribution, and adaptive frameworks that must be balanced against government accountability requirements.

Recommendations for Defense Acquisition Leadership

The following five recommendations are directed at Department of War acquisition leadership and are informed by the full body of evidence collected across all five research phases. Each recommendation is designed to be actionable within existing statutory authorities while advancing the institutional conditions necessary for OBC success.

Recommendation 1: Elevate Outcomes to the Requirements Stage

The FAR Part 11 Companion Guide reflects meaningful progress toward performance-oriented acquisition. However, by positioning OBCs as a contracting approach only, rather than as a requirements discipline, the guidance risks reinforcing a critical misunderstanding, namely that outcomes are a contract feature, rather than a strategic starting point. What distinguishes true OBCs from traditional PBCs is not contract form; it is the elevation of outcomes to the earliest stage of requirements development.

Part 11 should be revised to emphasize that defining outcomes is a fundamental requirements responsibility. By establishing outcomes at the requirements stage, defense organizations can ensure that acquisition planning, performance metrics, governance structures, and incentives are aligned from the start, making strategy the driver of contract design rather than an afterthought.

Recommendation 2: Reposition OBC Guidance in the FAR

Beyond the need to extend the outcome-based discussion to requirements discussed in recommendation 1, we find the existing placement of OBC guidance within FAR Part 37 (Service Contracting) creates a fundamental misconception about the scope and applicability of outcome-based contracting. This positioning inadvertently restricts practitioners’ understanding of OBCs as exclusively service-oriented mechanisms, when research evidence demonstrates their broader utility across diverse contract categories, including production and manufacturing contracts, construction and infrastructure projects, and major weapons system sustainment.

We recommend a two-part repositioning. First, the fundamental OBC definition should be established within FAR Part 2 (Definitions), which houses core acquisition terminology. This placement would establish OBCs as an enterprise-wide contracting concept with universal applicability across all defense acquisition activities. Second, the detailed description of when and how to structure OBCs should reside in FAR Part 16 (Types of Contracts), which is the natural home for contract structure and design guidance and where practitioners already look for vehicle selection decisions. This repositioning fundamentally reframes OBCs from a specialized services technique to a comprehensive contracting philosophy applicable across the complete spectrum of defense procurement.

Recommendation 3: Develop and Pilot OBC Governance Training

Our research consistently surfaces a critical workforce gap. Acquisition professionals lack the training, competency frameworks, and institutional support to develop and execute contract governance management plans, the very plans that make OBCs viable. The absence of governance capability is the single largest barrier to OBC success, and no amount of policy guidance will compensate for a workforce that has never been taught how to govern an outcomes-based relationship.

We recommend developing a dedicated OBC governance training curriculum that covers governance plan development, relationship management, joint performance review facilitation, adaptive decision-making, and outcome interpretation. Curriculum design should be grounded in role-based capability mapping, recognizing that contracting officers, program managers, requirements owners, and financial staff each face distinct governance gaps; the CCM Institute’s Contract Management Standard (CMS) 4 competency framework provides a useful foundation for structuring these differentiated learning paths. This training should be piloted across defense procurement organizations, beginning with agencies that have expressed interest in or have already attempted OBC approaches. Critically, the training must not be limited to contracting officers; it must reach program managers, contracting officer representatives, requirements owners, financial oversight staff, and industry partners who together form the governance ecosystem. Pilot results should be used to refine the curriculum and build the case for enterprise-wide adoption across the Department of War.

Recommendation 4: Develop a Portfolio Prioritization Schema for OBC Strategy

The current approach to OBC guidance reflects an understandable but limiting instinct: start with what is closest and most visible. The placement of OBC guidance in FAR Part 37 signals that the government's initial priority is professional services contracting. While professional services are a logical starting point, the conversation must be elevated to the entire defense acquisition portfolio.

We recommend that defense acquisition leadership, in partnership with GSA and other major procurement agencies, develop a structured prioritization schema that identifies which portfolio segments are most suitable for outcomes-based strategy enforcement and in what sequence. The schema should assess portfolio segments against the five critical success factors (Requirements, Data, Trust, Governance, Oversight) to determine readiness and potential impact. This effort must move beyond the reflexive focus on professional services to evaluate production contracts, IT modernization, facilities and construction, logistics, and major weapons system sustainment for outcomes-based strategy applicability.

Empirical research provides concrete guidance for this schema. Bohm et al. (2016) demonstrate that the optimal OBC form depends on contextual conditions: economic OBCs should be preferred for established, well-understood systems in technologically turbulent environments, while availability OBCs provide more appropriate risk boundaries for innovative or emerging technologies. Incorporating these contextual moderators alongside the five critical success factors would give defense leadership an empirically grounded basis for matching OBC type to portfolio segment characteristics.

Recommendation 5: Structure Contracts for Low-Risk OBC Piloting

One of the most significant barriers to OBC adoption in the defense acquisition context is the perceived risk of committing to an untested approach. We recommend implementing two contract structuring methods that leverage option periods to create safe, reversible OBC pilots.

Method A (Mid-Flight Conversion via Negotiated Option) applies to contracts already in performance. Teams add an option period to the end of the existing contract and negotiate a conversion of the nearest-term option to an OBC approach. The buyer and seller teams pilot the OBC model for that option term. If the pilot fails, both parties revert to the original contracting method for the remainder of the contract. The incentive for the contractor to participate at low or no additional cost is a guaranteed additional option period, a mechanism that rewards success if the OBC works well and provides a recovery runway if the OBC period proves difficult.

Method B (Parallel Option Tracks at Initial Award) applies to new contract awards. The government builds parallel option period structures into the contract from the outset. Each option period offers two tracks: a traditional or performance-based track and an OBC track. At the end of each performance year, the buyer and seller mutually agree on which track to exercise for the next period. Both methods lower the stakes of OBC experimentation and give practitioners a way to learn by doing, build governance muscle, and generate the performance data needed to assess OBC viability, all within existing contracting authorities and without betting the entire contract on an unfamiliar model.

Limitations

This study has several important limitations that should be acknowledged. The participant groups were broadly representative of federal acquisition organizations, spanning defense and civilian agencies across multiple functional roles and experience levels. The global expert interviews extended the evidence base to eight countries and multiple commercial sectors; however, the interview sample was limited in size (N = 14) and notably lacked gender balance, with 13 male and 1 female participant. This demographic skew reflects broader representation challenges in senior commercial and procurement leadership but limits the study’s ability to capture the full range of practitioner perspectives. Additionally, the research did not capture perspectives from state, local, or tribal procurement environments, where statutory frameworks, resource constraints, and contracting cultures differ meaningfully from the federal context. The findings should therefore be interpreted as well-grounded within the federal acquisition system and enriched by global commercial perspectives, but not necessarily generalizable to non-federal government contracting without further validation.

Federal policies governing OBCs, especially those found in supporting FAR materials, were still evolving while this research was being conducted. Some of the policy interpretations presented may change as guidance continues to develop. Since OBC implementation varies significantly from one agency to another, there is still insufficient empirical data from actual contract performance to conduct comprehensive long-term analysis. This uneven adoption pattern limits our ability to draw firm conclusions about performance trends over extended periods.

Future Research

This report represents the completion of Phase I in our research program, which encompassed five data collection phases integrating U.S. federal and global commercial perspectives. Phase II will pursue two primary objectives. The first is a comprehensive OBC case study analysis, documenting OBC applications across both government and private sector environments, spanning diverse contract categories (service delivery, production, and construction), examining cross-agency implementation, analyzing which of the five critical success factors were present or absent and how that influenced outcomes, and integrating international case studies for comparative benchmarking. The global expert interviews conducted during Phase I provide a foundation of international contacts and cross-cultural insights that will directly inform Phase II case selection and analysis.

The second objective is the development of a practical OBC decision framework: a structured decision-support tool that enables acquisition professionals to make informed determinations about OBC suitability and design. This framework will provide a systematic methodology for evaluating whether an OBC approach aligns with specific acquisition requirements, guidance for assessing the five critical success factors within particular procurement scenarios, and clear criteria to help practitioners distinguish between situations calling for outcomes-based requirements fulfilled through traditional vehicles versus situations warranting a specially structured OBC. This Phase II initiative will provide acquisition organizations with both the empirical evidence and practical tools necessary to make informed decisions about OBC implementation.

Conclusion

OBCs represent a major shift in how the Department of War approaches procurement. They move contracting away from detailed input requirements and process compliance toward measurable results and mission outcomes. This change is not simply procedural; it requires rethinking how contracts are designed, managed, and evaluated across the entire life cycle.

Our research shows that successful OBC implementation depends on five core areas: outcome-focused requirements, strong data capabilities, trust-based collaboration, effective governance structures, and accountability centered on results. Regulatory flexibility alone is not enough. Defense agencies must build workforce skills, strengthen cross-functional coordination, and redesign governance practices to support performance-driven relationships.

The theoretical foundation established in this report reinforces a fundamental point: contracts are not paperwork; they are value architecture. Every requirement, incentive, and performance metric is a statement about what the parties believe matters. When that statement faithfully represents the underlying purpose of the agreement, the contract fulfills its real function. When the document replaces the value narrative rather than serving it, even technically compliant contracts can fail to deliver intended outcomes (Finkenstadt, 2026).

The convergence of Graeber’s anthropological insight, Zeithaml’s perceived value framework, and Service-Dominant Logic yields a consistent message: value emerges through action, is co-created through interaction, is perceived rather than objective, and must be recognized through measurement to become actionable. The empirical evidence demonstrates that procurement professionals understand this intuitively; they value quality and outcomes over price by substantial margins when given meaningful signals about likely results (Finkenstadt, 2020; Graeber, 2001; Vargo & Lusch, 2004; Zeithaml, 1988). The practical imperative for defense acquisition leadership is a shift from buying tasks and things to buying outcomes and valuing them appropriately: defining requirements in terms of warfighter results, evaluating proposals based on outcome achievement potential, structuring contracts that align payment with results, and measuring what matters.

OBCs are not just a contracting technique. They are an institutional capability. To achieve measurable mission impact, the Department of War must invest in the governance, skills, and alignment necessary to turn outcome ambitions into sustained operational success.

Disclosures

Previous Reports

Portions of this report are based on a research project in print from the CCM Institute and the IBM Center for the Business of Government.

Use of Artificial Intelligence in Document Preparation

Portions of this document were prepared with the assistance of a large language model. Specifically, AI was used to synthesize the coauthors inputs as well as copyedit the document into the format required by ARP. All substantive analysis, research findings, policy recommendations, and intellectual contributions originate from the research team. The AI tool served as a drafting and synthesis aid; it did not generate original research, conduct independent analysis, or contribute novel conclusions. The human authors and research team reviewed, revised, and take full responsibility for all content presented in this report.

References

Bohm, E., Backhaus, C., Eggert, A., & Cummins, T. (2016). Understanding outcome-based contracts: Benefits and risks from the buyers' and sellers' perspective. *Journal of Strategic Contracting and Negotiation*.

CCM Institute. (2026). *AI in contracting 2026: From experimentation to impact*. Commerce & Contract Management Institute.

Federal Acquisition Regulation (FAR) companion guide, Parts 11, 16, and 37. (2025). U.S. Federal Acquisition Regulatory Council. https://www.acquisition.gov/content/far-overhaul-new-far-companion-guide

Federal Acquisition Regulation, 48 C.F.R. Parts 2, 11, 16, and 37 (2025).

Finkenstadt, D. J. (2020). *Essays on perceived service quality and perceived value in business-to-government knowledge-based services* [Doctoral dissertation, University of North Carolina at Chapel Hill].

Finkenstadt, D. J. (2026). *Contracts are not paperwork. They are value architecture* [Unpublished working paper].

Finkenstadt, D. J., Handfield, R., & Guinto, P. (2022). *How firms can plan for risk in a data saturated world: The goals, decisions, signals, data (GDSD) model*. California Management Review Insights.

Finkenstadt, D. J., & Zeithaml, V. A. (2020). *Perceived service quality in B2G knowledge-based services* [Working paper].

GAO. (2025). *F-35 Joint Strike Fighter: Actions needed to address late deliveries and improve future development* (GAO-25-107632).

Graeber, D. (2001). *Toward an anthropological theory of value: The false coin of our own dreams*. Palgrave Macmillan.

Guajardo, J. A., Cohen, M. A., Kim, S.-H., & Netessine, S. (2012). Impact of performance-based contracting on product reliability: An empirical analysis. *Management Science, 58*(5), 961–979.

National Contract Management Association (2025). Enhancing collaboration through effective relationship management. https://ccm.institute/research/publications/details/relationship-management

Tillipman, J. (2026). Abdicated judgment: AI tools and the future of reasoned decision-making in federal procurement. *Notice & Comment, Yale Journal on Regulation (Symposium on AI and the APA)*.

Vargo, S. L., & Lusch, R. F. (2004). Evolving to a new dominant logic for marketing. *Journal of Marketing, 68*(1), 1–17.

Vargo, S. L., & Lusch, R. F. (2008). Service-dominant logic: Continuing the evolution. *Journal of the Academy of Marketing Science, 36*(1), 1–10.

Zeithaml, V. A. (1988). Consumer perceptions of price, quality, and value: A means-end model and synthesis of evidence. *Journal of Marketing, 52*(3), 2–22.

The Value Hexagon of Warfighting Acquisition: Applying the Six Competing Constraints to Accelerate Warfighting Capability and Revitalize the U.S. Defense Industrial Base

Jeff Dunlap, CAPT USN (Ret.)—Lecture, Naval Postgraduate School. [Jeffrey.dunlap@nps.edu]

Abstract

This paper proposes the adoption of the Value Hexagon of Warfighting Acquisition. By adding three dimensions—Benefit/Value, Sponsor/Combatant Commander (CCMD) Priority, and Risk Tolerance—alongside the original three—Cost, Schedule, and Performance (C/S/P)—the Warfighting Acquisition System can shift from a risk-averse/risk-minimization focus to one of program value optimization and innovation.

The Value Hexagon is a visualization tool that provides Portfolio Acquisition Executives (PAE) with the analytical basis to make transparent decisions and swiftly, decisively shift funding within portfolios' authorized boundaries. Rather than shifting funds based on gut feel or political pressure, a PAE can show two Value Hexagon spider profiles side by side—a "Legacy Drifter" program bleeding resources and a "Rapid Responder" program starved of them—and document the trade; this is defensible to Congress and to the Secretary in a way that a narrative memo is not.

The PAE has the authority to ensure honest scoring on the Value Hexagon and to make real decisions about which programs continue to add warfighting value and which should be terminated. Requiring the decision authority to co-sign an explicit risk tolerance statement—rather than leaving risk as a hidden variable—creates accountability that does not currently exist in acquisition.

The Structural Flaw: The Compliance Culture

For decades, the Defense Acquisition System has relied on a Cost, Schedule, and Performance (C/S/P) framework to manage program health. The metrics collected indicated whether the acquisition programs were tracking the initial plan, but visibility into the value to the warfighter was based on static requirements that remained relatively stable in the post-Cold War era. This method is ill-suited for today's pacing challenge. Its core flaw is structural: by focusing exclusively on administrative compliance, it lacks a mechanism to assess whether a program actually delivers operational value. In reality, C/S/P are not independent of each other, and this oversimplified concept creates flawed expectations in managing Department of Defense (DoD) acquisition programs.

The Iron Triangle frames acquisition management as a zero-sum game: success means staying within baselined cost, schedule, and technical performance bounds. Program Managers (PMs) are therefore incentivized to treat any variance as a failure, regardless of whether that variance might have produced a better operational outcome. Earned Value Management (EVM) decomposes the program into a Work Breakdown Structure (WBS), and each WBS element gets its own cost and schedule baseline. The Iron Triangle compliance mentality is not just at the program level—it is replicated down to every work package. A team completing a sub-task is measured against its local cost and schedule targets, with no visibility into whether that sub-task is contributing to operational value. The compliance culture is therefore baked into the DNA of program execution at every level, not just the executive reviews (GAO, 2025; SAE International, 2019).

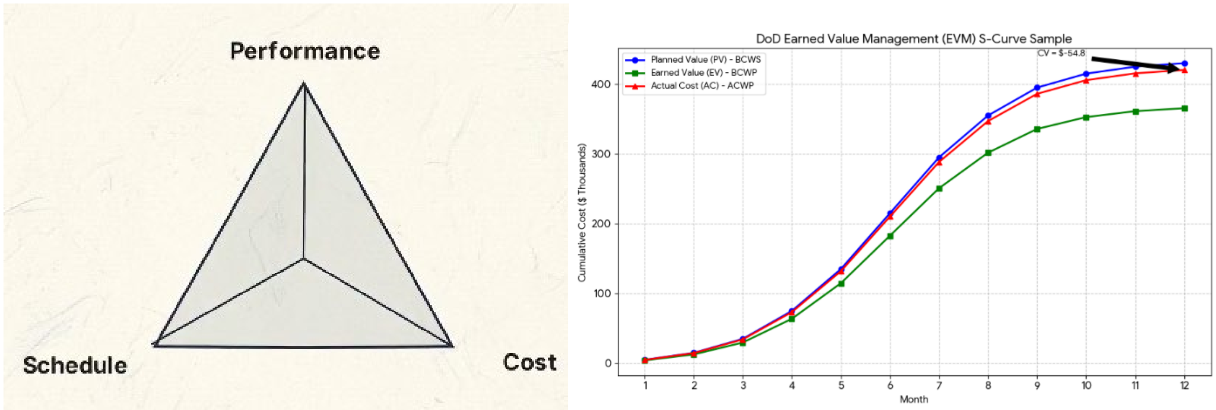

Figure 1. Concepts of the Iron Triangle and Earned Value Dimensions

The core EVM metrics—Cost Performance Index (CPI) and Schedule Performance Index (SPI)—tell you whether you are spending money at the rate you initially planned to spend it. A program that is burning budget on schedule looks "green" regardless of whether the work being done is producing anything useful. The fundamental measurement unit is planned value consumed, not capability delivered (SAE International, 2019).

The Core Point

The fundamental limitation of EVM is that compliance with cost and schedule baselines does not assure that the capability being produced remains operationally relevant to the warfighter.

Figure 2. Earned Value Reward Model and Impact on Warfighter

The deeper problem is that these models treat warfighter requirements as static artifacts. Under the legacy "Big A" Defense Acquisition System, requirements were often locked years—or even decades—before a system reached production. Because there was no formal "Value" axis, technical specifications inevitably drifted out of sync with a changing threat environment. The result was a compliance culture that reliably produces systems that satisfy outdated requirements while failing current operational needs. The Warfighting Acquisition System is the centerpiece of a major overhaul designed to shift the Pentagon from a slow, "requirements-based" process to a rapid, "solutions-based" model (initiated in late 2025 and codified in the 2026 National Defense Authorization Act [NDAA]). The strategy essentially treats acquisition as a warfighting function rather than a back-office administrative task, prioritizing "speed to capability" above traditional compliance (DoD, 2025).

The proposed Value Hexagon provides visibility, on an analytical basis, to enable swift, decisive decision-making. Authority without a documented rationale is politically and legally exposed. When a Portfolio Acquisition Executive (PAE) kills a program or redirects funding, the first question from Congress, IG, and GAO is “on what basis?” Under the current system, the honest answer is often “judgment”—which invites second-guessing, protest, and reversal. An auditable six-axis score that shows a “Legacy Drifter” consuming resources a “Rapid Responder” needs is a *defensible record*, not just a decision. The analytical basis is what converts authority into durable action.

> **The Core Argument**
> *Where existing frameworks measure fidelity to an administrative baseline, the proposed Value Hexagon measures fidelity to operational need. The distinction is not merely semantic—it determines what program managers are incentivized to optimize.*

Alignment with the Joint Force Requirements Process

The Value Hexagon is designed to operate in concert with the Joint Force Requirements Process (JFRP), which replaced the document-heavy Joint Capabilities Integration and Development System (JCIDS). Unlike JCIDS, which produced capabilities documents that took years to approve, the JFRP is a continuous cycle of Sponsor/ Combatant-command gap assessment (CJCS, 2026; GAO, 2025).

This cadence transforms the Value Hexagon from a static snapshot into a living program profile. When a threat shifts and a Sponsor/CCMD elevates a gap from Tier 3 to Tier 1, the program’s silhouette changes immediately, prompting a review of whether the cost, schedule, and risk settings need adjustment. In JFRP-era acquisitions, the high-priority/low-risk-tolerance mismatch will appear more frequently, not less, precisely because gaps are surfaced more quickly. The Value Hexagon makes those mismatches visible; the JFRP cadence provides natural decision points to correct them.

The Value Hexagon was designed to operate in precisely the JFRP environment. Every structural feature of the JFRP has a direct counterpart in the hexagon's six-axis framework, and the alignment is not coincidental—the hexagon's architecture was conceived specifically to make the JFRP’s operational logic visible at the program level.

Joint Operational Problems → Benefit/Value and CCMD Priority Axes

The most fundamental structural change in the JFRP is its pivot from requirements documents to problem statements. The JROC identifies and prioritizes Joint Operational Problems (JOPs) to communicate the Joint Force’s most pressing problems, focus analytical efforts, and align requirements to acquisition and resourcing.

JCIDS asked: what capability does the force need? The JFRP asks: what operational problem must the joint force solve? That reframing maps directly onto two hexagon axes working in tandem. This is a profound shift.

The ***Benefit/Value axis*** captures the answer to the JFRP question at the program level: what mission outcome does this capability deliver toward solving the identified JOP? Under JCIDS, benefit was implicit—assumed to be captured in the Key Performance Parameters. Under the JFRP, the JOP framing requires that benefits be stated explicitly in operational-effect terms, which is exactly what the Benefit/Value axis requires. A program that cannot articulate its contribution to a JROC-prioritized JOP cannot justify a high Benefit score, and that diagnostic failure is immediately visible in the hexagon silhouette.

The ***CCMD Priority axis*** captures the warfighter's tier assessment of the gap that the program addresses. The JROC oversees the prioritization of CCMD capability gaps and approval of joint requirements. That prioritization is the direct input to the CCMD Priority axis score. When the JROC elevates a gap from Tier 3 to Tier 1—because a JOP has risen in urgency—the hexagon profile changes immediately. The two axes together translate the JFRP's problem-centric logic into a program-level decision frame.

Capability Portfolio Management → Fiscal Sustainability Axis

The Requirements and Resourcing Alignment Board (RRAB) will align fiscal resources to JROC-prioritized JOPs during the Program Budget Review process. By exception, the RRAB may recommend modifying or terminating component activities that conflict with those priorities. This is the JFRP's portfolio management mechanism—the institutional instrument for ensuring that dollars follow operational problems rather than administrative inertia.

The hexagon's ***Fiscal Sustainability axis*** is the program-level expression of exactly this logic. It asks: Is this program consuming portfolio resources at a rate proportionate to its contribution to JROC-prioritized JOPs? A Legacy Drifter program that is administratively perfect but addresses a JOP the JROC has deprioritized scores poorly on Fiscal Sustainability—precisely the condition that should trigger RRAB attention. The hexagon makes that condition visible before the RRAB review cycle rather than surfacing it only when budget pressure forces the issue.

The Capability Portfolio Management (CPM) framework also introduces the concept of divestment as a legitimate portfolio outcome—recommendations will identify capabilities and constructs that require additional investment, maintenance of current investment levels, or divestment. The hexagon's Legacy Drifter and Gold-Plated Anchor archetypes are the diagnostic tool that makes divestment decisions auditable rather than political. The PAE can point to the silhouette—poor Benefit, low CCMD Priority, poor Fiscal Sustainability—as the documented analytical basis for a termination recommendation.

Continuous CCMD Gap Assessment → Dynamic Scoring Cadence

Under JCIDS, requirements were generated through a document-heavy process that could take years. Once a Capability Development Document was validated, the requirements were effectively frozen. The updated direction eliminates JROC validation of Service- and component-level requirements and refocuses the JROC on Joint Force Development, Joint Capability Integration, and CCMD requirements. By narrowing JROC oversight to Joint and CCMD priorities, the revised framework places greater emphasis on Service-level responsibility while strengthening alignment to joint warfighting outcomes and portfolio-based capability management.

The JFRP's continuous cadence transforms the Value Hexagon from a static scoring tool into a living program profile. JFRP's nominal timeline is 55 business days, representing a dramatically compressed cycle compared to JCIDS. Each JFRP cycle is a natural synchronization point at which the Value Hexagon's ***CCMD Priority*** axis must be rescored and the ***Fiscal Sustainability*** assessment updated to reflect whether the program's cost trajectory remains proportionate to its current JOP alignment.

This continuous cadence prevents the Value Hexagon from becoming the same kind of static artifact that JCIDS requirements documents have become. The JFRP provides the institutional clock that forces the Value Hexagon to be updated and provides the visualization that makes each JFRP update visible to decision authorities in a form they can act on immediately.

Joint Force Design → Schedule and Risk Tolerance Axes

The JFRP and JROC evaluate Joint Force Design (JFD), Joint Capability Integration (JCI), and CCMD requirements through a lens of JOPs underpinned by CPM, including Warfighting Acquisition System trade-space decisions. This is significant—the JFRP explicitly connects to the Warfighting Acquisition System (WAS) introduced by the 2026 NDAA, which reframes acquisition as a warfighting function rather than an administrative one.

The Value Hexagon's ***Schedule*** axis—reimagined as Time-to-Field rather than administrative milestone compliance—directly embodies the WAS's "speed to capability" mandate. The question is not whether the program is on its approved baseline, but whether the capability is arriving in time to address the JOP before the threat environment shifts again. The JFRP's 55-business-day cycle creates the operational urgency context within which that schedule judgment must be made.

The ***Risk Tolerance*** axis is where the Value Hexagon makes its most distinctive contribution to JFRP alignment. The JFRP's emphasis on speed and warfighting outcomes inherently demands that programs be authorized to accept risk. But the JFRP itself does not provide a mechanism for documenting the level of risk authorized for a given program. The Value Hexagon's Risk Tolerance Statement fills that gap—it converts the implicit risk posture embedded in every acquisition decision into an explicit, co-signed, auditable policy setting that can be compared against the CCMD Priority score at each JFRP refresh.

JROC Structure and the PAE Accountability Chain

The reoriented JROC is focused on JFD, JCI, and CCMD requirements. It conducts analytic due diligence to prioritize capability gaps, leveraging data-driven assessments to inform planning and budgetary decisions, and manages the JFRP to ensure its processes are streamlined, transparent, and focused on delivering capabilities at speed.

The Value Hexagon's governance design mirrors this structure at the program level. The JROC prioritizes JOPs and gaps at the joint level; the PAE translates those priorities into program-level trade-space decisions using the Value Hexagon as the analytical instrument. The hexagon does not replace the JFRP's deliberative process—it operationalizes its outputs at the milestone review level. When the JROC elevates a gap, the hexagon's CCMD Priority score rises, the risk tolerance mismatch diagnostic fires if the program's posture has not adjusted, and the PAE has an auditable basis for whatever decision follows.

The JFRP does not delineate actions required to satisfy acquisition rules and regulations. Services and Components must ensure their programs meet and satisfy the statutory and regulatory requirements for fielding capabilities to the warfighter. This is the precise gap the Value Hexagon fills. The JFRP tells the acquisition enterprise what the joint force needs and how urgently it needs it. The Value Hexagon tells program managers and decision authorities how to translate that signal into a defensible, continuously updated program posture—with the six axes providing the structure, the scoring rubric providing the rigor, and the spider chart providing the visualization that makes complex, multi-dimensional trade spaces accessible in a single portfolio review.

The Fundamental Alignment: Requirements Validation to Operational Problem-Solving

The deepest alignment between the Value Hexagon and the JFRP is philosophical. JCIDS was a validation system—it asked whether a stated requirement had been properly documented and approved. The JFRP is a problem-solving system—it asks whether the joint force's most pressing operational problems are being addressed with appropriate urgency and resources. The Value Hexagon is the measurement and visualization instrument that makes the problem-solving orientation operational at the program level.

Under JCIDS, a program manager asked, “Am I delivering against my validated requirements?” Under the JFRP/Value Hexagon framework, the question becomes: am I delivering against the joint force’s current operational problems at the speed and fiscal posture the situation demands? That shift—from requirements compliance to operational problem-solving—is what both the JFRP and the Value Hexagon are designed to institutionalize, and it is why they are structurally complementary rather than merely compatible.

The Value Hexagon: A Six-Dimension Framework

The Value Hexagon extends program management discipline within a six-dimensional framework that adds explicit measures of mission benefit/value, warfighter priority, and risk tolerance.

The Six Dimensions

The three retained legacy dimensions are reinterpreted as follows:

- **Cost:** Reimagined as **Fiscal Sustainability,** a portfolio-level affordability measure. The question is not whether the program is on its own baseline, but whether the portfolio can sustain it without crowding out higher-priority efforts. A program with high Benefit and CCMD Priority can still score poorly here if its cost trajectory consumes resources at a rate disproportionate to its JFRP tier.
- **Schedule:** Reimagined as “Time-to-Field.” Rather than a fixed administrative deadline, the schedule is balanced against the urgency of the identified capability gap.
- **Performance:** Technical Key Performance Parameters (KPPs) remain essential, but performance is now a tradeable variable—accepting an “80% solution” today to accelerate fielding is explicitly defensible.

The three new mission-driven dimensions are:

- **Benefit / Value:** The explicit mission outcome, measured in operational effects such as kill-chain compression (time to find, fix, and finish) or logistics footprint reduction.
- **Sponsor/CCMD Priority:** This axis ties a program’s importance to current campaign plans rather than decade-old documents.
- **Risk Tolerance:** An explicit policy setting, scored 1–10. This dimension transforms risk from a hidden program variable into a deliberate management choice.

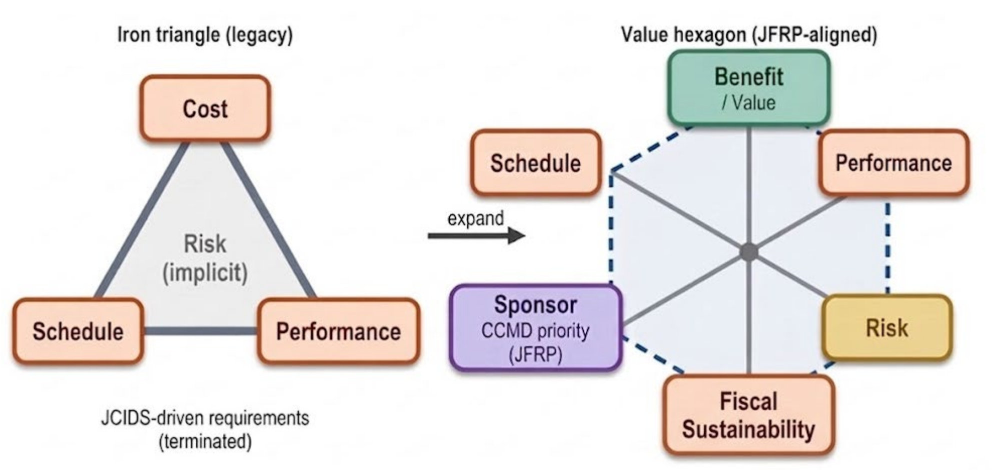

Figure 3. JCIDS to JFRP Value Hexagon Management Focus

Table 1. Comparison of the Iron Triangle and Value Hexagon Dimensions

Dimension	Iron Triangle	Value Hexagon
Cost / Fiscal Sustainability	Implicit; program-level baseline compliance, no portfolio visibility.	Portfolio affordability axis; scored against crowding risk to adjacent programs.
Schedule	Fixed milestones; delays are failures.	Time-to-Field; balanced against urgency of the capability gap.
Performance	Rigid KPP thresholds; hard to trade.	Variable; tradeable against benefit for faster delivery.
Benefit / Value	Implicit; assumed to be captured in specs.	Apex axis; measured in operational effects (e.g., kill-chain compression).
Sponsor/CCMD Priority	Static; derived from years-old requirements.	Dynamic (Tiers 1–3); refreshed each JFRP cycle.
Risk Tolerance	Hidden; managed as a shadow of cost/schedule.	Explicit policy setting; a deliberate PAE choice.

Risk Tolerance

The Risk Tolerance axis works differently from the other five. The other dimensions ask "how well is the program performing?"—higher is better. Risk Tolerance asks, "How much variance has the PAE authorized this program to absorb?" It is a policy setting, not a performance grade. A score of 3 does not indicate a troubled program; it indicates a program where the PAE has deliberately chosen a conservative posture, which may be entirely appropriate for a low-urgency effort.

The critical diagnostic insight is the relationship between Sponsor/CCMD Priority and Risk Tolerance. When these two scores are mismatched—High Priority paired with low Risk Tolerance—the Value Hexagon surfaces a structural tension that the Iron Triangle would have left invisible. A program the warfighter urgently needs cannot be acquired at the speed of relevance if the acquisition organization has no authority to accept meaningful risk.

> **Key Insight: The Mismatch Problem**
> *When Sponsor/CCMD Priority scores are high, but Risk Tolerance scores are low, the resulting Value Hexagon profile reveals a structural tension that conventional C/S/P reporting cannot surface. This mismatch requires explicit resolution by the PAE—either through adjusting the authorized risk posture, restructuring the acquisition pathway, or formally acknowledging timeline constraints.*

Three Readings of the Mismatch Profile

When a program shows High Priority alongside low Risk Tolerance, there are three distinct root causes—each with a different resolution:

- **Resourcing:** Low risk tolerance often reflects insufficient funding margin, immature technology, or an inexperienced program office. If the gap is genuinely Tier 1, this posture must be corrected by adding risk margin or escalating to the PAE.
- **Strategy:** The low-risk posture may be deliberate: a performance threshold set below operational need to buy time. The correct response is a spiral or incremental acquisition strategy—field a lower-capability increment quickly, then iterate at the next JFRP refresh (DoD, 2020).
- **Culture:** If multiple programs in the portfolio show this profile, it may indicate systemic risk aversion—scores suppressed to protect programs from scrutiny. The Value Hexagon is only as useful as the honesty of the scoring; this pattern warrants an audit.

All three paths converge on the same obligation: the PAE must resolve the tension explicitly and on the record, not allow it to dissolve into milestone paperwork.

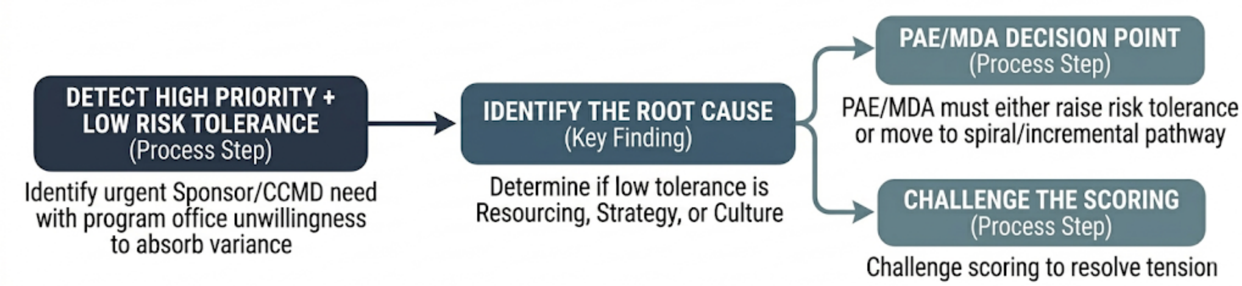

Figure 4. Tension in the Value Hexagon Drives Action

Program Archetypes

The true utility of the Value Hexagon lies in its analytically based visualization. Rather than parsing lengthy status reports, decision authorities can immediately read a program's "silhouette"—the shape formed when its six scores are plotted—and identify mismatches between intent and execution.

> **Key Insight: The Shape of the Spider Chart is Important**
> *The shape of the spider chart—not merely its area—communicates the program's strategic posture. Large, balanced silhouettes indicate programs delivering broad value across all dimensions, while elongated or compressed profiles signal specific mismatches requiring PAE attention.*

The Value Hexagon Spider Chart

The Value Hexagon spider chart is a single visual that tells you the strategic story of an acquisition program in about ten seconds. Instead of reading a 40-page program status report or parsing rows of earned value data, a PAE or Milestone Decision Authority (MDA) looks at one shape and immediately sees where the program is strong, where it is strained, and whether the tensions within it demand action.

It is called a "spider chart" because the six axes radiate outward from a center point like spokes on a wheel. Each axis is scored on a scale of 1 to 10. The program's scores are plotted along those spokes and connected, forming a closed polygon—the program's "silhouette." The shape of that silhouette is the message.

The Six Axes

Each axis measures something genuinely distinct. None of them can be derived from the others, which is what makes the six-axis structure honest.

Axis	The Question It Answers
Benefit / Value	What mission outcome does this capability actually deliver?
Performance	Is the system meeting its technical specifications?
Fiscal Sustainability	Can the portfolio absorb this program without crowding out higher-priority efforts?
CCMD Priority (JFRP)	How urgently does a Combatant Command need this, right now?
Risk Tolerance	How much variance has the PAE/MDA authorized this program to absorb?
Schedule	Is the capability arriving in time to matter?

The first three axes on the left side of the Value Hexagon—Benefit, Performance, and Fiscal Sustainability—answer the question “what are we getting and at what cost to the portfolio?” The three on the right—CCMD Priority, Risk Tolerance, and Schedule—answer the question: “How urgently do we need it and how are we set up to deliver it?” A healthy program is balanced across both sides. A Full scoring rubric is provided in the Appendix.

How to Read the Shape

This is the core executive skill. The silhouette communicates faster than any number.

- A large, balanced hexagon—scores of 7 or above on all six axes—signals a program that is delivering high mission value on time, at a fiscally sustainable pace, with a clear warfighter sponsor, and with appropriate risk authorization. These programs deserve protection from budget pressure, not scrutiny.
- A small, compressed hexagon—low scores across most or all axes—signals a program that has lost its strategic rationale. It may be executing perfectly against an outdated requirement, but it is not delivering meaningful value to the joint force. This is the Legacy Drifter archetype: administratively perfect, operationally irrelevant.
- An elongated or lopsided shape—high on some axes, low on others—is where the real intelligence lies. Elongated shapes signal structural tension. Two specific mismatches demand immediate executive attention:
    - High Sponsor/CCMD Priority + Low Risk Tolerance: The joint force urgently needs this capability, but the program has not been authorized to move fast. This is the Urgency Trap. The silhouette is thin and stretched. The PAE must either raise the authorized risk tolerance or honestly acknowledge that the timeline cannot be met—because the current posture demands speed, it has not enabled it.

- High Performance + Low Benefit/Value + Low Fiscal Sustainability: The program has over-engineered itself into irrelevance. It meets extreme technical specifications that drive costs up, but those specifications no longer translate into operational advantage. This is the Gold-Plated Anchor. The shape is heavy on the Performance spoke and thin everywhere that matters to the warfighter.

Why Shape Matters More Than Area

A common mistake when first reading spider charts is to focus only on the polygon's total area—bigger is better. That instinct is wrong, and understanding why is important.

A program can have a large total area by scoring high on Performance and Fiscal Sustainability while scoring low on Sponsor/CCMD Priority and Benefit. That large area tells you the program is well-managed and cheap—but it says nothing about whether anyone in the joint force actually needs it. The shape reveals which dimensions are driving the area, and that is what determines whether the investment is justified.

The PAE should train their eye to ask three questions when looking at any silhouette:

1. Are the Benefit and Sponsor/CCMD Priority spokes pointing outward? If not, the program lacks a strategic rationale regardless of how well it is executed.
2. Is the Risk Tolerance spoke consistent with the CCMD Priority spoke? If Priority is high and Risk Tolerance is low, there is a structural problem that no amount of good program management can resolve.
3. Is Fiscal Sustainability proportionate to Benefit and Priority? A program scoring 9 on Priority and 3 on Fiscal Sustainability is consuming Tier 1 resources faster than its contribution justifies—it is cannibalizing the portfolio.

Figure 5. Sample Archetypes (Urgency Trap, Legacy Drifter, Rapid Responder)

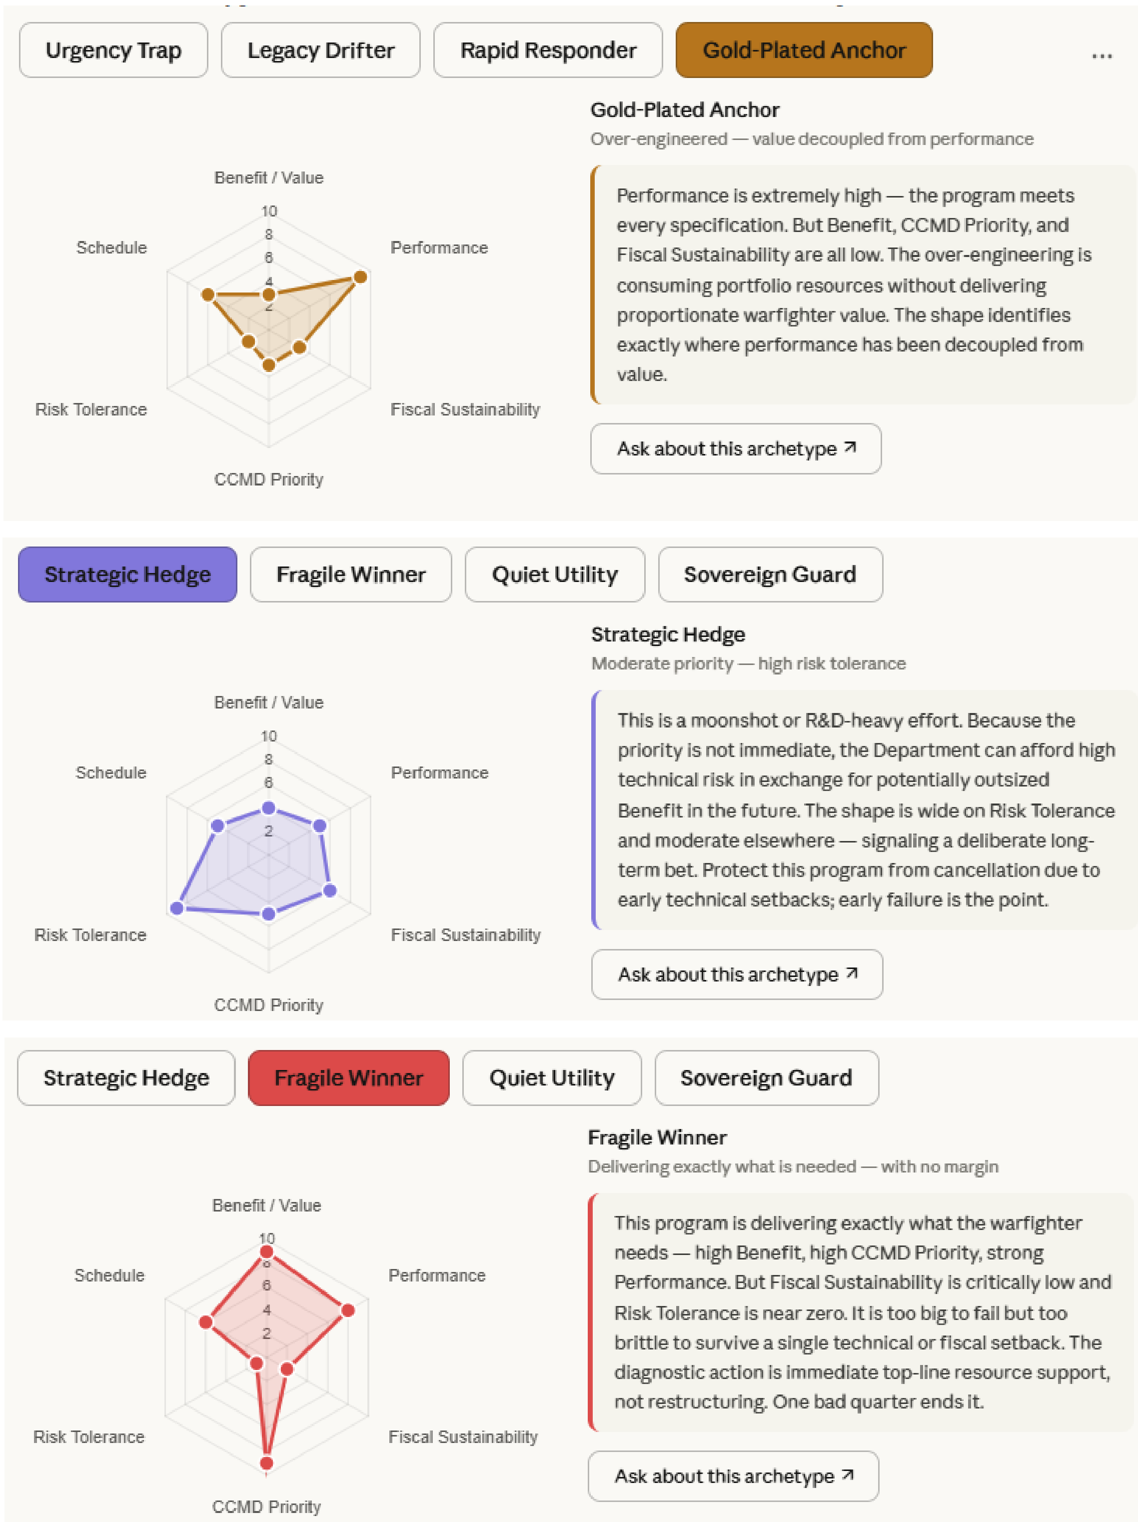

Figure 6. Sample Archetypes (Gold Plated, Strategic Hedge, Fragile Winner)

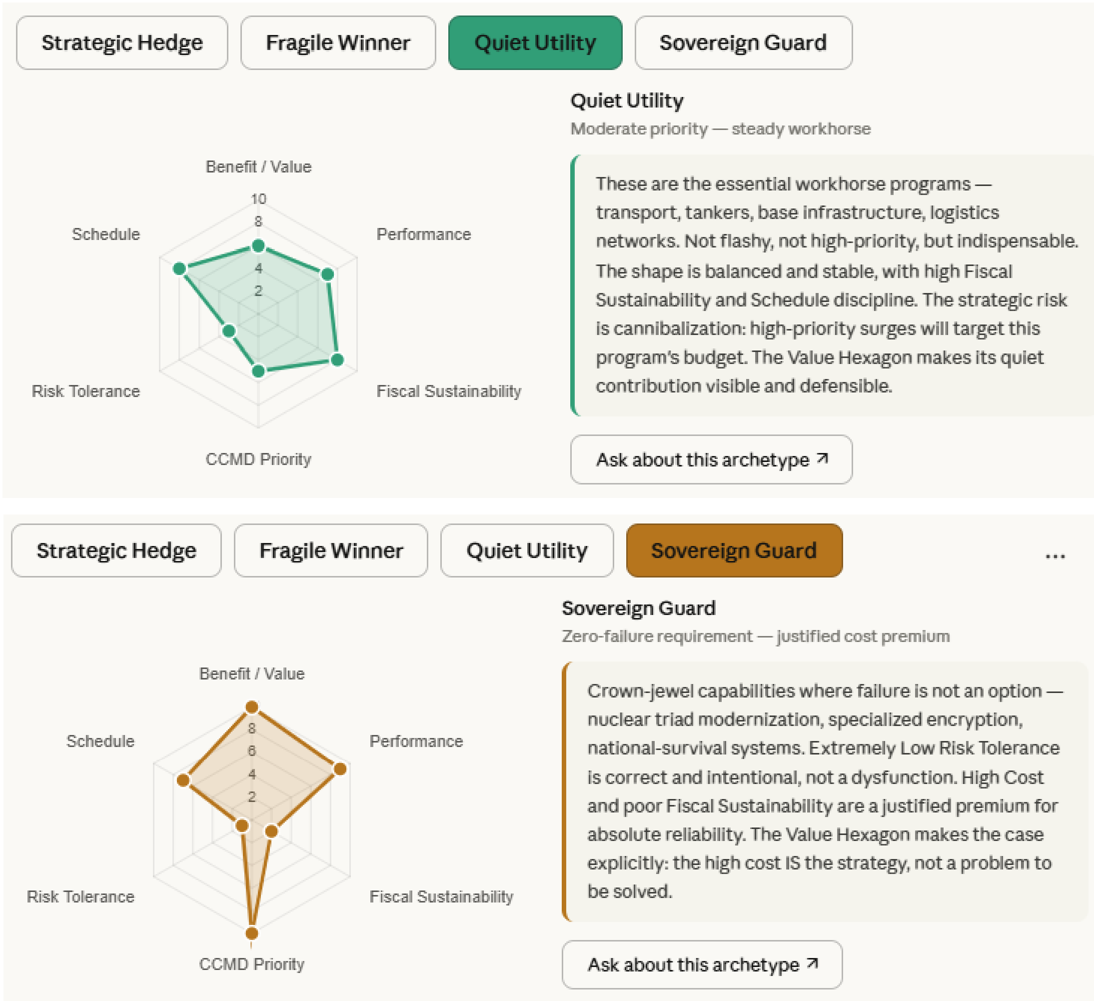

Figure 7. Sample Archetypes (Quiet Utility, Sovereign Guard)

Comparing Programs Side by Side

The spider chart's most powerful use is portfolio comparison. When two programs are plotted side by side—or overlaid on the same chart—the PAE can support transparent, auditable, and defensible funding reallocation decisions for Congress, the Inspector General, and the Secretary.

Consider this example. A *Legacy Drifter* scores 8 on Performance, 7 on Fiscal Sustainability, and 8 on Schedule—it is executing beautifully. But it scores 2 on CCMD Priority, 1 on Benefit, and 3 on Risk Tolerance. Its silhouette is heavy on the administrative compliance spokes and nearly absent on the mission value spokes.

Placed next to a *Rapid Responder*—which scores 9 on CCMD Priority, 8 on Benefit, 8 on Risk Tolerance, but 5 on Performance and 4 on Fiscal Sustainability because it is being starved of resources—the case for reallocation is visual, not narrative. The PAE does not have to argue from judgment. The two shapes make the argument.

Operationalizing the Hexagon Across the Life Cycle

Implementing the framework requires process changes at every acquisition phase, not merely new charts.

Requirements Phase

The JFRP gap submission replaces the former Initial Capabilities Document as the program's value baseline. Mission-outcome metrics—not KPPs alone—are established before development begins and tied explicitly to the Sponsor/CCMD's tiered priority.

Analysis of Alternatives

Alternatives are scored across all six dimensions with explicit weighting visible to the PAE. A more expensive option that addresses a Tier 1 JFRP gap with higher mission benefit can be selected transparently, with the trade documented rather than obscured.

Portfolio Management

Across the portfolio, the Value Hexagon enables direct comparisons between programs competing for limited resources. A program delivering high benefit against a Tier 1 gap with manageable risk may warrant continued investment even if it shows cost or schedule variance—a trade the Iron Triangle could not represent and the Value Hexagon makes both explicit and defensible.

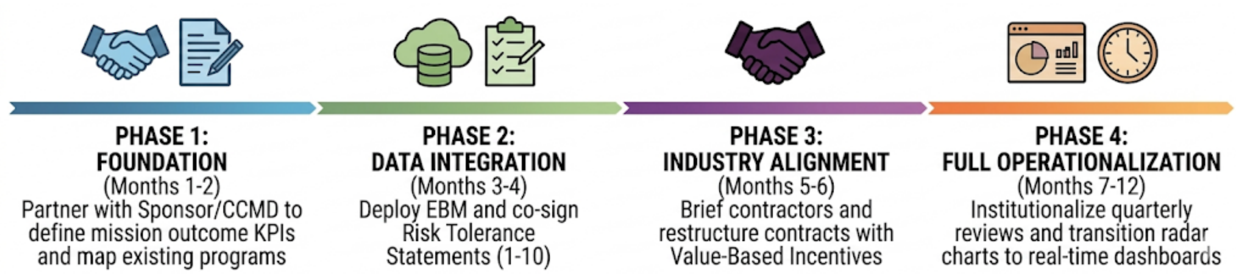

Figure 8. PAE Implementation Roadmap

Phase 1: Foundation (Months 1–2)

Objective: Establish the initial Value Hexagon profile for the program.

- Partner with the Sponsor/CCMD to translate existing requirements into 3–5 mission-outcome KPIs (e.g., "Reduction in target identification time" rather than "processor speed").
- Conduct a independent baseline assessment to assign the program's current archetype.
- Update the Program Management Plan to include the Value Hexagon spider chart as the primary reporting artifact for milestone reviews.

Phase 2: Metrics Integration (Months 3–4)

Objective: Make scoring auditable and objective.

- The PAE and PM co-sign a Risk Tolerance Statement (scored 1–10) that defines observable conditions for each score level.
- Establish a direct data link to the Sponsor/CCMD's JFRP gap-submission cycle for real-time Priority axis updates.

Phase 3: Industry and Stakeholder Alignment (Months 5–6)

Objective: Signal the new demand standard to the industrial base.

- Brief prime contractors and key suppliers on how performance will now be judged against "Mission Benefit," not just cost and schedule compliance.

- Incorporate value-based incentive structures into upcoming contract actions (e.g., bonus for reducing kill-chain time by an additional 15%).
- Conduct a mock milestone review using Value Hexagon archetypes to train senior leaders on reading program silhouettes.

Phase 4: Full Operationalization (Months 7–12)

Objective: Institutionalize the model as the single source of program truth.

- Run quarterly portfolio silhouette reviews to identify systemic patterns (e.g., inadvertent accumulation of Gold-Plated Anchors).
- Migrate Spider charts from manual slides to automated dashboards drawing from the Common Data Environment.
- Conduct an annual Value Audit comparing benefit realized against Risk accepted to calibrate the next JFRP cycle.

Critical Success Factors

Table 2. PAE Implementation Critical Success Factors

Factor	Requirement
Honest Scoring	Axes must be scored on observable, auditable conditions—not set to “green” for political safety.
PAE Backing	The PAE must protect the PM when authorized risk is taken and fails, provided it was aligned to a Tier 1 Sponsor/CCMD Priority.
Sponsor/CCMD Access	The program office must maintain a direct line to the warfighter’s operational planners to keep the Benefit axis current.
Dynamic Baselines	The program must be prepared to pivot its archetype when the JFRP signals a shift in threat priority.

Revitalizing the Defense Industrial Base

The Paradigm Shift: From Compliance to Mission Optimization

The U.S. Defense Acquisition landscape is undergoing its most significant structural evolution in fifty years. The transition to the Joint Force Requirements Process (JFRP) and its diagnostic centerpiece, the Value Hexagon, represents a shift from administrative compliance to mission optimization. For the Defense Industrial Base (DIB), this is not merely a policy change; it is a competitive imperative. Executives must now weaponize mission-value data to justify higher-margin, high-benefit solutions that provide the “speed of relevance” required for modern deterrence. The following table identifies the structural flaws of the legacy model and the advantages of the JFRP-aligned framework:

For industry leaders, the “So What?” is transformative. The Value Hexagon effectively “bridges the Valley of Death” for non-traditional contractors by making a prototype’s mission value visible to senior leadership long before a formal Program of Record is established. Traditional primes, conversely, must shift from “green-washing” underperforming programs to proving "Earned Benefit." Understanding the six dimensions of this hexagon is the prerequisite for portfolio realignment and capturing market share in the JFRP era.

The Six Dimensions of the Value Hexagon

The Value Hexagon is a holistic decision-making framework designed to anchor program health to mission-driven axes. It allows the DoD to move beyond checking boxes and to maximize deterrence by adjusting resources based on the evolving realities of the battlefield.

Fiscal Sustainability

- *Industry Impact:* Executives must demonstrate that their program's cost trajectory is proportionate to its JFRP priority tier. A program consuming Tier 1 portfolio resources while addressing a Tier 3 gap is fiscally unsustainable regardless of its technical performance. Conversely, a high-cost program addressing a Tier 1 gap with no alternative scores well here.

Schedule

- *Industry Impact:* Speed is a primary differentiator. Firms that can deliver "80% solutions" today will outperform those promising perfection in a decade.

Performance

- *Industry Impact:* Performance is now a tradeable variable. Contractors should be prepared to trade exquisite specs for faster delivery or higher mission utility.

Benefit / Value

- *Industry Impact:* This is the apex metric. Every engineering decision must be mapped to a warfighting effect rather than a static document.

CCMD Priority

- *Industry Impact:* Contractors must track CCMD needs "left" of the RFP. A high-tier signal indicates the DoD's willingness to authorize increased risk or funding.

Risk Tolerance

- *Industry Impact:* High tolerance signals an opportunity for aggressive, "fail-fast" prototyping. Low tolerance demands reliability-maximized engineering.

Analytical Note: Risk Tolerance as a Policy Lever

Note on Fiscal Sustainability: Unlike the old "Cost" axis—which measured whether a program was tracking its own approved baseline—Fiscal Sustainability is a portfolio-level measure. It asks whether the program is consuming resources at a rate proportional to its JFRP priority tier and whether it is crowding out adjacent programs that may have higher operational urgency. A program can deliver excellent mission value and still score poorly on Fiscal Sustainability if its cost trajectory is structurally unsustainable within the portfolio.

Fiscal sustainability is genuinely independent—a program can score high on both Benefit and Priority and still be fiscally unsustainable. Unlike the other five axes, Risk Tolerance is a policy setting, not a performance grade. A low score (e.g., 3) is not a "bad" grade; it reflects a deliberate choice by the Milestone Decision Authority (MDA) to hold the line tightly.

Strategic Portfolio Mapping: Analyzing Program Archetypes

"Silhouette analysis" using spider charts allows leadership to identify mismatches between intent and execution at a glance. By mapping portfolios to these archetypes, firms can identify which programs to scale, pivot, or divest.

The Capture Strategy: Winning in the Value-Based Environment

Winning in the JFRP era requires shifting business development "left" to influence requirements through CCMD intimacy before the RFP is released.

Directives for Proposal Innovation:

1. Pitch the Silhouette: Use spider charts in white papers to visually demonstrate how your solution maps to the government's desired archetype (e.g., trading extreme specs for a 40% improvement in Time-to-Field).
2. Propose Value-Based Contracts: Move away from "lowest price" mindsets. Propose incentive fees tied to bonuses for specific "kill-chain" reduction metrics.
3. Frame Risk Transparency: For high-priority gaps, explicitly define the technical risks required to meet the timeline and the specific "Risk Tolerance" setting required from the MDA.

Resolving the "Urgency Trap": A Diagnostic Flowchart

When a contractor identifies a high-priority/low-risk mismatch, they must guide the PAE through a root-cause analysis:

- Resourcing (Insufficient Margin): If the low risk tolerance stems from immature technology or thin funding, the contractor must advocate for a raised risk tolerance through added margin or TRL increases.
- Strategy (Deliberate Constraint): If the constraint is intentional, the contractor should propose a Spiral/Increment approach—using Other Transaction Authorities (OTAs) to field a lower capability fast and iterate.
- Culture (Systemic Aversion): If risk scores are being suppressed to "protect" the program from scrutiny, the contractor must challenge the score and request a re-baselining against the actual JFRP gap tier.

The Value Hexagon Functions as a Demand Signal to the Private Sector

The sustained emphasis on administrative compliance metrics has created structural incentives within the DIB that favor expertise in program management and regulatory navigation over speed-to-innovation. This is a rational response to the demand signals the acquisition system has historically transmitted, rather than a failure of industrial capacity per se. The Value Hexagon corrects this by making mission benefit an explicit, scoreable, and contractable variable (DoD, 2022; Schwartz et al., 2024).

Five Mechanisms for DIB Revitalization

- **Clearer demand signals:** Scoring programs across all six dimensions tells industry exactly which combinations of attributes DoD values, allowing firms to shape R&D and bids toward high-priority gaps rather than generic "better, cheaper, faster" targets.
- **Innovation incentives:** When a more expensive but higher-benefit solution can win a Tier 1 JFRP gap, firms are rewarded for innovation and speed-to-field rather than margin compression. This attracts venture-backed defense-tech entrants who cannot compete on administrative compliance alone.
- **Transparent source selections:** Explicit, auditable six-axis scoring reduces perceived arbitrariness in award decisions, lowering protest risk and business volatility—which in turn unlocks private capital investment in defense-oriented firms.

- **Incremental production opportunities:** Hexagon profiles showing high priority but low risk tolerance naturally point to spiral acquisition strategies, creating recurring production and upgrade contracts rather than single large programs of record.
- **Portfolio stability for critical suppliers:** The hexagon enables the DoD to sustain investment in programs delivering high mission value even under cost or schedule pressure, stabilizing demand for specialized industrial capabilities (advanced munitions, space, C2) that are vulnerable to cancellation during short-term budget cycles.

Conclusion

The Iron Triangle was a useful tool for an era of stable requirements, predictable threats, and a primary concern with fiscal discipline. None of those conditions reliably exists today, and a framework designed to optimize for compliance cannot be repurposed to optimize for deterrence.

The Value Hexagon does not discard what the Iron Triangle established. Cost discipline, schedule performance, and technical rigor remain embedded in three of the six axes. What the hexagon adds is the structure to ask the question the Iron Triangle was never designed to answer: is this capability worth what we are spending, and does the joint force actually need it right now? Anchored to the JFRP's continuous combatant-command-driven prioritization cycle, the framework ensures that answer is updated as the threat evolves rather than frozen at the moment a requirements document was signed.

The practical result is a portfolio management instrument that converts authority into durable action. When a PAE reallocates funding from a Legacy Drifter to a Rapid Responder, the hexagon provides the auditable six-axis record that survives Congressional scrutiny, Inspector General review, and GAO inquiry—not because it is bureaucratically airtight, but because it is operationally honest. The shape of the spider chart is defensible precisely because it reflects the CCMD commander's gap assessment, the MDA's co-signed risk posture, and the value actually delivered—not a program office's narrative.

The deeper purpose, however, is not administrative. It is cultural. The acquisition system has spent decades measuring what is easy to measure—dollars spent, milestones hit, specifications met—and calling that program health. The Value Hexagon measures what the system was always supposed to deliver: operational capability at the speed of relevance. Adoption of this framework, paired with the institutional conditions necessary to sustain honest scoring, is not merely an acquisition reform. It is a step toward an institution that earns the trust of the warfighter it serves.

Appendix: Notional Scoring Rubric

Value Hexagon — Scoring Rubric (1-10)

Score	Benefit / Value	Performance	Fiscal Sustainability
1	No measurable mission outcome. Nice-to-have; no campaign plan dependency.	KPPs not met; system fails basic functions in operational environment.	Program is fiscally unsustainable. Cost trajectory is accelerating; EAC growth is crowding out Tier 1 and Tier 2 programs in the portfolio. Immediate portfolio-level intervention required.
2	Marginal improvement to existing capability; not linked to any CCMD gap.	Fewer than half of KPPs met; significant gaps vs. threshold requirements.	Severe sustainability risk. Unit cost growth or TOA consumption is disproportionate to the program's JFRP tier. Portfolio is absorbing Tier 3 costs at Tier 1 rates.
3	Supports a lower-tier gap; benefit qualitative and hard to quantify.	Most KPPs at threshold only; no objectives achieved; marginal utility.	Poor sustainability. EAC trending significantly above original estimate; funding profile unstable across the FYDP. Adjacent programs are at risk of underfunding as a result.
4	Validated gap addressed; benefit measurable but modest (minor sortie improvement).	All KPP thresholds met; some objectives achieved; adequate in limited conditions.	Below-average sustainability. Cost trajectory shows upward pressure; funding stability requires attention before next milestone. Portfolio rebalancing may be needed at JFRP refresh.
5	Meaningful improvement to one mission area; supports a Tier 3 JFRP gap.	All thresholds met; majority of KPO objectives achieved; reliable in primary cases.	Moderate sustainability. Program is consuming its authorized share of portfolio resources with minor variance. Stable but offers no margin against unexpected cost growth.
6	Significant improvement across multiple mission areas; supports Tier 2 JFRP gap.	All thresholds and most objectives met; performs well across expected conditions.	Acceptable sustainability. Cost trajectory is flat or gently declining. Portfolio share is proportionate to JFRP priority tier. No immediate portfolio risk.
7	Addresses Tier 1 JFRP gap; quantified improvement to kill-chain or logistics.	All KPPs at or above objective values; exceeds baseline in primary mission.	Good sustainability. Unit costs stable or declining; EAC at or below original estimate. Program is not crowding adjacent efforts and has adequate funding margin.
8	Critical to a specific campaign plan; removal would degrade CCMD concept.	Objectives exceeded across all KPPs; validated in live operational testing.	Strong sustainability. Demonstrated cost discipline over multiple reporting periods. Portfolio obligation share is declining relative to capability contribution.
9	Joint-force-wide impact; enables/protects strategic objective in priority theater.	Significantly exceeds all objectives; validated under adversarial/degraded conditions.	Excellent sustainability. Unit cost reduction confirmed through production or competition. Program is freeing resources that can be redirected to higher-priority gaps.
10	Existential mission criticality; no substitute; absence creates unacceptable NDS risk.	Best-in-class; sets new baseline for capability area; no known shortfalls.	Exceptional sustainability. Net contributor to portfolio health: declining unit costs, stable funding profile, zero crowding of adjacent programs, surplus available for Tier 1

Value Hexagon — Scoring Rubric (1-10)

Score	CCMD Priority (JFRP)	Risk (Tolerance)	Schedule
1	Not referenced in any CCMD gap submission; service-internal requirement only.	Zero tolerance; any variance unacceptable; program rigidly constrained.	IOC delayed >2 years against APB; recovery plan not credible.
2	Referenced in CCMD submission but not prioritized; informational mention only.	Extremely low tolerance; minor variance requires immediate escalation; no margin.	Delay of 1-2 years; SPI <0.8; critical path at risk; recovery under development.
3	Addresses a Tier 3 JFRP gap for a single CCMD; limited joint applicability.	Low tolerance; small management reserve; near-zero technical risk required.	Delay of 6-12 months; SPI 0.8-0.9; recovery needs resources or descope.
4	Tier 3 gap across multiple CCMDs; broader but still lower-priority need.	Limited tolerance; modest reserve; accepts low risk on mature technologies.	Delay of 3-6 months; manageable with current resources; no IOC breach.
5	Tier 2 JFRP gap for a single CCMD; meaningful operational priority.	Moderate tolerance; adequate reserve; accepts some schedule slip for capability.	Minor slip 1-3 months; SPI >=0.95; within normal variance; no APB breach.
6	Tier 2 gap across multiple CCMDs; joint force relevance clear and documented.	Balanced tolerance; accepts cost/schedule variance if performance is preserved.	On schedule against APB; SPI >=1.0; all major milestones tracking green.
7	Tier 1 JFRP gap for one CCMD; directly linked to a priority campaign plan.	Elevated tolerance; significant reserve; accepts risk on higher-TRL innovations.	Ahead of schedule 1-3 months; SPI >1.0; milestone compression possible.
8	Tier 1 gap across multiple CCMDs; cross-theater relevance validated in JFRP cycle.	High tolerance; flexibility explicitly approved; willing to prototype and fail fast.	Tracking 3-6 months early; demonstrated schedule discipline across phases.
9	Top-tier joint force priority; cited in Joint Warfighting Concept guidance.	Very high tolerance; structured for experimentation; failure is a planned outcome.	Significantly ahead; early IOC achievable; delivering capability ahead of plan.
10	Highest-priority gap in JFRP cycle; supports NDS pacing threat; SECDEF visibility.	Maximum tolerance; OTA/rapid acquisition vehicle; speed over cost certainty.	Exceptional performance; IOC substantially early; new benchmark for speed.

1-3 = Unsustainable / High Portfolio Risk (red) | 4-6 = Moderate Sustainability (yellow) | 7-10 = Strong / Portfolio-Enhar

References

CJCS. (2026). *CJCSI 5123.01J: Charter of the Joint Requirements Oversight Council (JROC) and implementation of the Joint Force Requirements Process (JFRP)*. U.S. Joint Chiefs of Staff. https://www.jcs.mil/Portals/36/Documents/Library/Instructions/CJCSI%205123.01J.pdf

Defense Acquisition University. (2022). Valley of death. Defense Acquisition University glossary. https://govmates.com/defining-the-valley-of-death/

Department of Defense, Office of the Under Secretary of Defense for Research and Engineering. (2023). *National defense science & technology strategy 2023*. https://www.cto.mil/wp-content/uploads/2024/05/2023-NDSTS.pdf

DoD. (2020, January 23). *DoD Instruction 5000.02: Operation of the adaptive acquisition framework* (incorporating change 1, June 8, 2022). Office of the Under Secretary of Defense for Acquisition and Sustainment. https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500002p.PDF

DoD. (2022, February). *State of competition within the defense industrial base*. Office of the Under Secretary of Defense for Acquisition and Sustainment. https://media.defense.gov/2022/Feb/15/2002939087/-1/-1/1/STATE-OF-COMPETITION-WITHIN-THE-DEFENSE-INDUSTRIAL-BASE.PDF

DoD. (2025, November 7). *Warfighting Acquisition System implementation directive establishing Portfolio Acquisition Executives*. Office of the Secretary 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

GAO. (2025, June). *Defense acquisition reform: Persistent challenges require new iterative approaches to delivering capabilities with speed* (GAO-25-108528). https://www.gao.gov/products/gao-25-108528

National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119-60, 140 Stat. 1 (2026). [Warfighting Acquisition System provisions]

SAE International. (2019). *Earned value management systems* (EIA-748D). SAE International. https://www.sae.org/standards/eia748d-earned-value-management-systems/

Schwartz, M., Peters, H., Guo, A., & Siripurapu, A. (2024). *The U.S. defense industrial base: Background and issues for Congress* (CRS Report No. R47751). Congressional Research Service. https://crsreports.congress.gov/product/pdf/R/R47751

Portfolio Management Competency Standards

Dr. Robert F. Mortlock, COL, USA (Ret.)—is a professor of the practice and associate chair for acquisition sciences for the Department of Acquisition, Finance and Manpower at the Naval Postgraduate School in Monterey, CA. He holds a PhD in chemical engineering from the University of California, Berkeley, an MBA from Webster University, an MS in national resource strategy from ICAF, and a BS in chemical engineering from Lehigh University. [rfmortlo@nps.edu]

Michael S. Bobulinski, CDR, USN (Ret.)—is currently serving as the Embedded Training Readiness Project Manager at NSWC Dahlgren Division–Dam Neck Activity. He holds an MS in program management and an MS in computing technology both from the Naval Postgraduate School and a BS in aerospace engineering from Pennsylvania State University. [michael.s.bobulinski.civ@us.navy.mil]

Abstract

Department of Defense (DoD) acquisition programs and professionals have been under scrutiny for years. Direction has been provided, over time, to adopt civilian program management practices within the DoD. The Project Management Institute (PMI) sets and manages civilian project, program, and portfolio management standards and certifications. This study assesses DoD alignment with PMI standards and focuses on portfolio management competency standards. In both this study and previous research, gap analysis methodology (both qualitative and quantitative approaches) was applied. This research found a nearly 60% alignment between the DoD program management (PM) competency model and industry standards. The research recommends an alignment between DoD standards and the industry-accepted PMI standards for portfolio management.

Keywords: portfolio management, program management, gap analysis, NDAA acquisition guidance, acquisition reform and innovation

Introduction

The National Defense Authorization Act (NDAA) for Fiscal Year (FY) 2021 established portfolio management as a management process for the acquisition of defense weapons systems to reduce costs, improve the optimization of the investments in weapons systems, and increase acquisitions efficiency (NDAA, 2021; Office of Congressman Wilson, 2020). Furthermore, the FY 2020 NDAA established the certification requirements for the members of the acquisition workforce based upon third-party accredited, internationally recognized standards (NDAA, 2020). This is a shift from the current strategy with program management (PM) focus within defense acquisition and necessitates study of the alignment between existing PM competency standards with portfolio management (PfM) competency standards. The FY2021 NDAA provides portfolio management guidance: “(b) Portfolio Management – The Secretary of Defense shall establish capabilities for robust, effective, and data-driven portfolio management described in subsection (a)(1)(C)” (NDAA, 2021).

Despite direction from the FY2021 NDAA for full implementation to be completed by 2023 (NDAA, 2021), the Department of Defense (DoD) has failed to do so. In contrast, the Project Management Institute (PMI) introduced a portfolio management certification in 2014, the Portfolio Management Professional (PfMP) certification (PMO Advisory, n.d.). Government Accountability Office (GAO) reports over the years have been critical of the DoD’s project, program, and portfolio management. Some excerpts from several GAO reports follow:

- The Department of Defense (DoD) is not effectively using portfolio management to optimize its weapon system investments. (Sullivan, 2015, Highlights page)

- DoD does not have a policy to guide portfolio management across the department that fully reflects key best practices. The policy is also not current and DoD is not implementing it. (Sullivan, 2015, p. 15)
- In nearly all cases, the military services could improve their practices by learning from ideas and initiatives being used . . . by commercial companies and ensuring that civilian and military personnel have similar opportunities to develop. (Sullivan, 2018, p. 26)
- DoD partially concurred with our 2015 recommendation related to improving portfolio management. . . . However, DoD has yet to fully address the recommendation. (Oakley, 2021, p. 66)
- GAO has long reported on needed improvements to DoD's portfolio management practices. (Mak, 2022, Highlights page)

The Defense Acquisition Workforce Improvement Act (DAWIA) of 1990 requires the DoD to establish a professional acquisition workforce, and the DoD has focused on this requirement since 1992 (Gates et al., 2024, p. 1). In FY2021, the DoD consolidated 14 career fields into seven functional career fields; however, portfolio management was not recognized as a distinct career field (Gates et al., 2024, p. 2). An Acquisition Innovation Research Center study from the University of Maryland in September 2023 on Portfolio Performance Analysis and Visualization went as far as stating, "DoD is not following industry standards for portfolio management" (Driessnack & Johnson, 2023, p. 10, figure 3).

This study assesses the alignment of the DoD's general PM competency standards with the PMI's competency standards for portfolio management. A previous study from the Naval Postgraduate School (NPS), "Gap Analysis of Department of Defense Program Management Competency Standards in Preparation for the Shift to Portfolio Management in Defense Acquisitions," from December 2021 found that the DoD was roughly 41% aligned with industry standards. The authors of that study used gap analysis as their approach for determining alignment (Stewart et al., 2021). This study follows a similar research analysis methodology.

Primary Research Question: Is there alignment between DoD acquisition workforce PM competencies and PMI portfolio management professional standards?

The scope of this research was narrowed to the analysis of the competency standards required for acquisition professionals and the potential application of new standards to encompass PfM. Structural, budgetary, statutory, and design implications that require further research may materialize when implementing the shift from program-centric management to PfM. The shift in focus from program to portfolio management is a significant endeavor for the DoD that requires analysis of existing competency standards to determine the applicability of the existing standards and the requirement for developing new standards. Applying nationally accepted industry standards to PfM competencies in the DoD may be a vital component to improving the acquisition system.

Background and Literature Review

Understanding the definitions and basis for competency standards is critical to analyzing the DoD's alignment with industry standards. As previously mentioned, the GAO has been critical of the DoD and most federal agencies and organizations on their implementation of project, program, and now portfolio management. DoD Weapon System Acquisition has been on the GAO High-Risk List since 1990 and continues to struggle (Oakley, 2025).

Currently, acquisition career fields established by the DAWIA and managed by the Services' Directors of Acquisition Career Management (DACMs) do not formally recognize portfolio manager as a career field separate and distinct from a program manager, creating a potential gap between the competency standards and the requirement for PfM. While the Section 809 Panel, the Office of Management and Budget (OMB), and the GAO have advocated for PfM for years, change has been challenging to adopt (Ahern & Driessnack, 2019; Section 809 Panel, 2018a, 2018b, 2019a, 2019b, 2019c; Sullivan, 2015; Thompson & Johnson, 2019). In the private industry sector, an organization's shift from program-centric acquisition strategies to PfM strategies stems from two drivers: the need to make rational investment decisions that deliver organizational benefits and the need to optimize resources to ensure the efficient delivery of those benefits (Young & Conboy, 2013). PfM achieves these benefits by pooling resources and analyzing how decisions made about one product affect the other products in the portfolio and portfolio priorities writ large. Additionally, the defense acquisition enterprise comprises numerous commands with their own goals, agendas, and interpretations of policies (GAO, 2020). These organizations change leaders and priorities every three or four years. This "fragmented adhocracy" makes implementing change difficult (Young & Conboy, 2013, p. 1090). Finally, implementing PfM requires competent professionals. According to Young and Conboy (2013), competence is "the ability to do something well" (p. 1091). PfM requires a common competency standard as the metric to train and evaluate acquisition professionals. Identifying gaps in the competency standards assists in updating and codifying a standard that can be used as a common thread to synchronize PfM efforts across the defense acquisition enterprise.

Within the DoD, significant challenges exist preventing the full implementation of PfM. One reason for the absence of standards related to PfM is a lack of clarity. In the academic community and industry, there has been confusion as to what constitutes PfM. The term often gets used interchangeably with PM, project management, and multi-project management (Young & Conboy, 2013). This same confusion about what constitutes PfM exists within the DoD. In the private sector, corporate PfM practices and procedures have been undervalued and under-researched, leading to an identified gap between the direction and means available to implement PfM (Gutiérrez & Magnusson, 2014; Heising, 2012; Kock et al., 2020; Li et al., 2014; Petit, 2012). Despite many medium and large corporations applying PfM principles and tools to make strategic decisions, "academic research has not kept up with the realities and needs of the corporate world" (Nippa et al., 2011, p. 64). The lack of corporate PfM-focused research, combined with the NDAA statutory push to leverage PfM based on industry accepted standards within the DoD, presents a need to conduct focused PfM research to recognize its value.

PfM is an approach that commercial companies use to optimize investments (Sullivan, 2015). It starts with understanding customers' needs and desires and then prioritizing acquisition opportunities while accounting for resource constraints. Once the opportunities are prioritized, business cases are created, reviewed, and "assessed against others in the portfolio" (Sullivan, 2015, p. 5). Resources, established criteria, competing products, and the organization's strategic goals are all considered during the assessment. This process continues "until only those alternatives with the greatest potential to succeed" are added to the portfolio (Sullivan, 2015, p. 5). Therefore, the DoD would only create new programs through a holistic portfolio analysis process (Sullivan, 2015).

A PfM strategy improves the defense acquisition procedures in three significant ways. First, it requires acquisition professionals to assess investments collectively at the enterprise and component levels rather than as independent initiatives at the Service level. Second, it uses "an integrated approach to prioritize needs and allocate resources" to align with strategic goals (Sullivan, 2015, p. 7). Last, it empowers leaders to make investment decisions and provides a mechanism to hold them accountable for the outcome (Section 809 Panel, 2019a).

Current defense acquisition procedures measure success through cost, schedule, and performance metrics for individual programs with acquisition program baselines. However, these measures do not allow program managers to develop optimal solutions across a range of capabilities and customer needs. Therefore, at times, they can be detrimental to the larger, strategic mission. Additionally, they provide little insight into the value the program offers to the customer. Lastly, they do not allow flexibility because they incentivize stability and avoid new requirements. Instead, PfM should be judged on things such as “customer satisfaction, user acceptance or reject rates, user productivity improvements, mission effectiveness enhancements, and many others that relate to value and return on investment” (Shultz, 2020, p. 47). Additionally, there must be a mechanism to measure the success of things such as rapid prototyping. These may include metrics such as “time to deliver knowledge points, cycle time to build virtual prototypes, number of failures and lessons learned, and time to mature prototypes into fieldable capabilities” (Shultz, 2020, p. 47).

Defining PfM is of particular importance in the DoD because the terms *program*, *portfolio*, and *project* are often used interchangeably by defense acquisition professionals at all levels. The PMI defines a portfolio as “a collection of projects, programs, subsidiary portfolios, and operations managed as a group to achieve strategic objectives” (PMI, 2017b, p. 6). Figure 1 and Table 1 show the relationship and comparison between the definitions accepted for projects, programs, and portfolios and between the disciplines of project, program, and portfolio management.

Table 1. Comparative Overview of Portfolio Program, and Project Management (PMI, 2017b, p. 6)

Organizational Project Management			
	Projects	**Programs**	**Portfolios**
Definition	A project is a temporary endeavor undertaken to create a unique product, service, or result.	A program is a group of related projects, subsidiary programs, and program activities that are managed in a coordinated manner to obtain benefits not available from managing them individually.	A portfolio is a collection of projects, programs, subsidiary portfolios, and operations managed as a group to achieve strategic objectives.
Scope	Projects have defined objectives. Scope is progressively elaborated throughout the project life cycle.	Programs have a scope that encompasses the scopes of its program components. Programs produce benefits to an organization by ensuring that the outputs and outcomes of program components are delivered in a coordinated and complementary manner.	Portfolios have an organizational scope that changes with the strategic objectives of the organization.
Change	Project managers expect change and implement processes to keep change managed and controlled.	Programs are managed in a manner that accepts and adapts to change as necessary to optimize the delivery of benefits as the program's components deliver outcomes and/or outputs.	Portfolio managers continuously monitor changes in the broader internal and external environments.
Planning	Project managers progressively elaborate high-level information into detailed plans throughout the project life cycle.	Programs are managed using high-level plans that track the interdependencies and progress of program components. Program plans are also used to guide planning at the component level.	Portfolio managers create and maintain necessary processes and communication relative to the aggregate portfolio.
Management	Project managers manage the project team to meet the project objectives.	Programs are managed by program managers who ensure that program benefits are delivered as expected, by coordinating the activities of a program's components.	Portfolio managers may manage or coordinate portfolio management staff, or program and project staff that may have reporting responsibilities into the aggregate portfolio.
Monitoring	Project managers monitor and control the work of producing the products, services, or results that the project was undertaken to produce.	Program managers monitor the progress of program components to ensure the overall goals, schedules, budget, and benefits of the program will be met.	Portfolio managers monitor strategic changes and aggregate resource allocation, performance results, and risk of the portfolio.
Success	Success is measured by product and project quality, timeliness, budget compliance, and degree of customer satisfaction.	A program's success is measured by the program's ability to deliver its intended benefits to an organization, and by the program's efficiency and effectiveness in delivering those benefits.	Success is measured in terms of the aggregate investment performance and benefit realization of the portfolio.

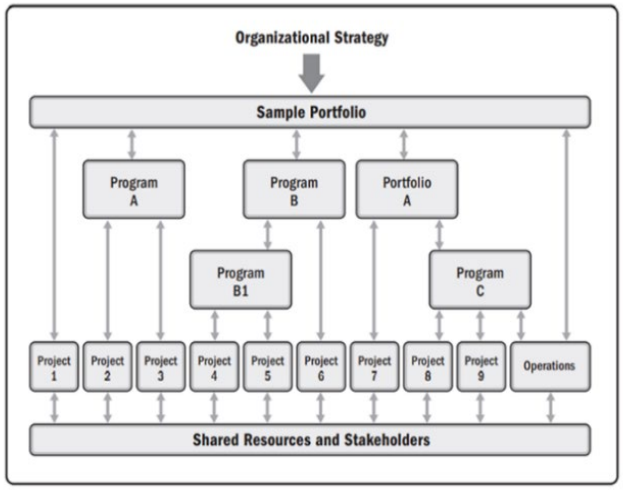

Figure 1. Portfolios, Programs, and Projects: High-Level View. (PMI, 2017b, p. 4)

A portfolio is a way to hedge against risk by pooling resources. Hence, a portfolio should be made with a clear strategy and priorities that the manager can use to make decisions. If portfolio managers are given a set of missions or capabilities they must meet, they can then analyze the assets, products, and programs within the portfolio available to fulfill that mission. The manager can then identify gaps in the portfolio where the DoD must allocate resources. These gaps inform how funding, personnel, and R&D should be allocated, all while keeping within the overarching strategy of the portfolio. Portfolio managers are not overly invested in the success or failure of any project or program but instead focus on how individual programs perform holistically within the portfolio (PMI, 2017b). Success is determined based on "aggregate investment performance and benefits realization of the portfolio" (PMI, 2017b, p. 6).

Within defense acquisition, portfolio management has technically been required since 2008 with the establishment of DoD Directive 7045.20, *Capability Portfolio Management*, and the framework for portfolio management has been in place since the establishment of Program Executives Officers (PEOs) in the 1990s (Office of the Under Secretary of Defense for Acquisition and Sustainment [OUSD(A&S)], 2023). However, "no substantial changes to the program approach have materialized," as the majority of projects maintained the program-centric model because the overall structure of the defense acquisition system "is not well suited for portfolio-based management" (Section 809 Panel, 2019a, p. 77). Despite the creation of PEOs in the 1990s and the direction for portfolio management, "PEOs were not assigned any additional duties in statute or DoDD 5000.01 to accomplish portfolio management. . . . Instead, they are midlevel managers," without being responsible or held accountable for a portfolio management baseline (Section 809 Panel, 2019a, p. 77).

Over the last several decades, the U.S. government sponsored numerous efforts, studies, panels, and reports regarding the requirement for defense acquisition to undergo significant reform, depart from the historical PM approach, and manage acquisitions in a portfolio-centric model (Anton et al., 2019; Biedenbach & Müller, 2012; Distel et al., 2020; Ewing et al., 2013). These efforts were codified by the Section 809 Panel on Streamlining and Codifying Acquisition Regulations as established by the direction contained in the FY2016 NDAA. The purpose of the Section 809 Panel was to "review the acquisition regulations . . . with a view toward streamlining and improving the efficiency and effectiveness of the Defense acquisition process" (Section 809 Panel, 2017, p. 5). The Section 809 Panel "identified portfolio management as a priority for reform, recommending not only a change in investment processes but a shift away from the decades-old program-centric acquisition model" (Shultz, 2020, p. 44). Specifically, the Section 809 Panel's (2019a) Recommendation 38 is to "implement best practices for portfolio management" (p .17) and includes the following language:

> Moving defense acquisition from a highly centralized, program-centric model with stovepipe-driven requirements, budget, and acquisition processes to a collaborative, decentralized, portfolio-centric framework entails nothing more than implementing management best practices. The move would yield timely, flexible, agile, cost-effective, and technologically innovative weapon systems acquisition and sustainment. Portfolio management is no longer in its infancy; there are standards and best practices that [the] DOD can use while implementing the recommended multitiered capability portfolio framework. (Section 809 Panel, 2019a, p. 84)

While some acquisition professionals argue that PfM already occurs due to the previous instructions and directives, "each program navigates the acquisition life cycle independently [and] programs design, develop, test, and produce individual systems that meet a defined set of requirements within an allocated budget" (Janiga & Modigliani, 2014, p. 13) regardless of classification under a portfolio.

According to DoD Instruction 5000.66 *Defense Acquisition Workforce Education, Training, and Career Development Program*, a competency is a "measurable pattern of knowledge, skills, abilities, behaviors, and other characteristics that an individual needs to perform work roles or occupational functions successfully. Competencies are used to develop acquisition training and education standards" (OUSD[A&S], 2022). DoDD 7045.20 *Capability Portfolio Management* establishes the policy for using capability portfolio management (CPM) across the DoD. CPM is defined as "a disciplined management approach to align, prioritize, and optimize investments, requirements, risks, resources, research, and developments around a set of capabilities to achieve a set of mission objectives" (OUSD[A&S], 2023).

In December 2024, the acting assistant secretary of defense (ASD[A]) signed a memorandum titled "Program Management Functional Career Field Competencies" (Office of the Assistant Secretary of Defense [ASD], 2024). The ASD(A) memorandum details functional competency units with a listing of competencies within each unit (shown in Table 2). It is significant to note that this memorandum references DoDI 5000.66 but does not reference DoDD 7045.20. It is equally significant to point out that the memorandum does not address portfolio management, nor does it establish competencies specifically for portfolio management. Despite these disconnects, the competencies areas, units of competencies, and competencies in Table 2 form the basis for this research as the assumed DoD PM standard to compare to industry-accepted standards.

Table 2. Functional Competencies.
(OASD, 2024)

Management Competency Units and Competencies December 1, 2024

Acquisition Management	Business Management	Technical Management
Capability Integration Planning	**Contract Management**	**Engineering Management**
Requirements Management	Market Research	Technical Planning
Acquisition Program Strategic Planning	Pre-Solicitation Planning and Execution	Requirements Decomposition
Business Case Development	Source Selection & Negotiations	Decision Analysis
Acquisition Law and Policy	Contract Administration	Configuration Management
Acquisition Policy and Best Practices	Contracting Approaches	Digital Engineering
Contractual Laws, Regulations, and Obligations	**Financial Management**	**Digital Literacy**
Financial Mgmt Laws, Directives, and Policies	Financial Planning	Machine Learning
Stakeholder Management	Programming	Artificial Intelligence
Political Savvy	Budget Formulation	Software Acquisition
External Situational Awareness	Budget Execution	**Test and Evaluation Mgmt**
Program Execution	Cost estimates	Test Planning: Preparation, Integration, Analysis Reporting
Risk/Opportunity Management	**Business Acumen**	**Product Support Mgmt**
Teaming	Internal/External Politics	Product Support Planning
Program Oversight	Financial Terms, Motivations, Incentives	Product Support Management
Resource Management	Public/Private Industry Differences	Supply Chain Mgmt and Supply Chain Risk Mgmt
Technology Management	Challenges/Constraints & Competitive Environment	Diminishing Manufacturing Sources & Materiel Shortages
Program Planning	Business Capture	
Pathway Selection		
Tailoring Acquisition Approach		
Executive Leadership		
Foundational Competencies	**Leading Change**	**Results Driven**
Interpersonal Skills	Creativity & Innovation	Accountability
Integrity / Honesty	Vision	Decisiveness
Communicate Effectively	Flexibility & Resilience	Customer Service
Continual Learning	**Leading People**	Problem Solving
Public Service Motivation	Conflict Management	**Building Coalitions**
Technical Credibility	Developing Others	Influencing / Negotiating
Digital Literacy	Team Building	Partnering

For industry standards, the International Standards Organization (ISO) and American National Standards Institute (ANSI) provide for clarity on definitions of competencies and standards. The PMI publishes ANSI-accredited standards for project, program and portfolio management used as the industry standards (PMI, n.d.). The ISO is a global organization that defines and publishes standards across industries (International Standards Organization [ISO], n.d.-a), and ANSI is the United States's member body in the ISO. ANSI provides oversight and accredits standards within the United States across American industries to include the U.S. government (ANSI, n.d.-a.). ISO 9000 defines competence as the "ability to apply knowledge and skills to achieve intended results" (ISO, n.d.-b, section 3.10.4). ISO views a standard to be "a document established by consensus and approved by a recognized body that provides rules, guidelines, or characteristics for activities or their results, aiming for the optimum degree of order in a given context" (American Society for Quality, n.d.). Similarly, ANSI defines a standard as "a document that provides requirements, specifications, guidelines or characteristics that can be used consistently to ensure that materials, products, processes, and services are fit for their purpose" (ANSI, n.d.-b., What is a Standard?).

ANSI recognizes the PMI as the leading independent organization in developing the standards and certifications for program, project, and portfolio management (Karnes, 2020). The PMI Project Management Professional (PMP), Program Management Professional (PgMP), and Portfolio Management Professional (PfMP) credentials are widely accepted and recognized internationally throughout industry to demonstrate an individual's commitment to meeting the highest levels of professionalism. The PfMP certification is one of the most rigorous offered and requires an extensive amount of experience. PfMP applicants must have a minimum of eight years of professional business experience and four to seven years of unique, nonoverlapping professional portfolio management experience (PMI, 2017a). This does not mean that the applicant must be the senior portfolio manager but, instead, must just have worked in an organization that uses the PfM construct. *The Standard for Portfolio Management,* 4th edition, explains various tasks related to the six recognized performance domains shown in Figure 2 (PMI, 2017b).

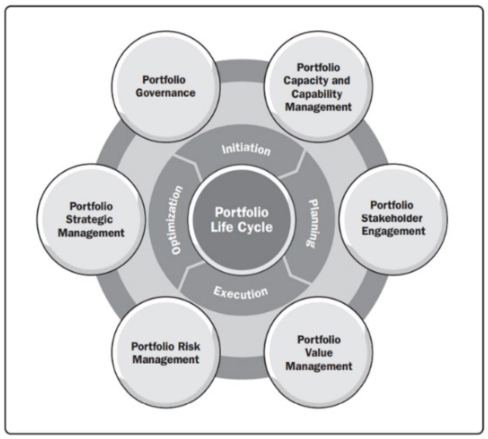

Figure 2. Portfolio Management Performance Domains. (PMI, 2017b, p. 10).

Data Sources and Methodology

This research incorporated mixed quantitative and qualitative methods that resulted in a competency gap analysis by mapping the alignment of the current DoD Program Management Functional Career Field Competencies (OASD, 2024) to the PMI (2013) *PfMP Examination Content Outline* domains and tasks, which aligns to *The Standard for Portfolio Management,* 4th edition (PMI, 2017b).

Shown in Table 2, the DoD Program Management Functional Career Field Competencies consist of four areas, including Acquisition Management (AM), Business Management (BM), Technical Management (TM), and Executive Leadership (EL; OASD, 2024). Within the four areas, there are 17 units of competency that contain 62 different competencies (OASD, 2024). The PMI (2013) *PfMP Examination Content Outline* served as the primary data source for industry PfM competency standards. The PMI designed the PfMP exam to reflect the required skills of portfolio management professionals (PMI, 2013). The PfMP exam "measures and evaluates appropriately the specific knowledge and skills required to function as a portfolio management professional" (PMI, 2013, p. 1). The purpose of the exam is to ensure that each required element of PfM is accurately measured to validate competency. This purpose aligns with the goal of the *DoD PM Career Field Competency Standards*. The exam outline lists five domains and weighs each in terms of importance for assessment. This weight is depicted by the percentage of questions on the exam, as outlined in Table 3. The five assessed domains are Strategic Alignment, Governance, Portfolio Performance, Portfolio Risk Management, and

Communications Management. Each of these domains includes subordinate tasks. The Appendix provides a detailed explanation of the tasks within the portfolio management domains. This research assumed that PMI's domains and DoD's functional areas were equivalent in nature, and that PMI's tasks within their domains and DoD's competencies and sub-competencies were equivalent in nature.

Table 3. Portfolio Management Professional Examination Domains and Weights. (PMI, 2013, p. 3).

Domain	Percentage of Items on Exam
Strategic Alignment	25%
Governance	20%
Portfolio Performance	25%
Portfolio Risk Management	15%
Communications Management	15%

A lexicographic analysis of keywords and the principal purpose of each DoD PM competency was matched to PMI PfMP domain tasks. The researchers created a competency alignment matrix with three classifications of alignment: No Discernible Alignment, Partial Alignment, or Full Alignment. The assessment of alignment was based on the following criteria:

- No Discernible Alignment indicated that no current DoD PM competency standard fit the description of a PMI-stated task.
- Partial Alignment indicated that one or more keywords or the general purpose of the DoD PM competency or sub-competencies related to the PMI stated task.
- Full Alignment indicated that an existing DoD PM competency standard matched the PMI stated task to the degree that included several exact word matches or clearly aligned descriptions, purposes, or applications.

To assess a quantitative measure of alignment, an Alignment Score scale was defined:

- No Discernible Alignment = 0
- Partial Alignment = 0.5
- Full Alignment = 1

Each PMI PfMP task was assigned an alignment score based on the qualitative assessment. Within each PfMP domain, the average score was calculated (i.e., the total score of all tasks divided by the total number of tasks within the domain). The average scores indicate the degree to which the DoD is already postured to transition to train and assess portfolio management skills based on its current PM competency standards.

After reviewing and matching all applicable DoD PM competency standards to the PMI domains and tasks, barriers to implementation (BTI) were identified. A shift from a PM-centric to a PfM-centric strategy will inherently require policy and organizational changes. The assessed barriers signal to defense acquisition decision-makers the areas where the researchers perceive that implementation would be the most challenging. The BTI approach used to analyze alignment included the following:

- No BTI as practices that already occur within the DoD
- Low BTI as changes that the DoD could implement immediately with little to no change in personnel structure or additional policy concerns

- Medium BTI as changes that would require either significant changes in policy or personnel structure
- High BTI as changes that would require both significant personnel structure and policy changes

To assess a quantitative measure of BTI, the following BTI rating scale was defined:

- No BTI = 0
- Low BTI = 1
- Medium BTI = 2
- High BTI = 3

Each PMI PfMP task was assigned a BTI rating using this scale based on the qualitative assessment. Within each PfMP domain, the average score was calculated (i.e., the total score of all tasks divided by the total number of tasks within the domain). This rating indicated the assessed degree of difficulty in implementing portfolio management standards based on current DoD acquisition practices, personnel, and policy.

Results

Figure 3 depicts the alignment between the PMI PfMP competency standards and the DoD competency standards broken down by PfMP domain. The overall average alignment of the two standards was nearly 60%. However, within each domain, those alignment scores vary significantly. In the domains of Strategic Alignment and Governance, the DoD is less than 50% aligned with PfMP standards, while in the domain of Communications Management, the two standards are 100% aligned. When evaluating the overall alignment score, it is important to recognize the weights of each domain from the *PfMP Examination Content Outline* (PMI, 2013). The most heavily weighted domains (Strategic Alignment and Portfolio Performance) exhibit 44% and 65% alignment percentages, respectively. This is significant because the weights from the exam represent the importance of the domain in evaluating competency.

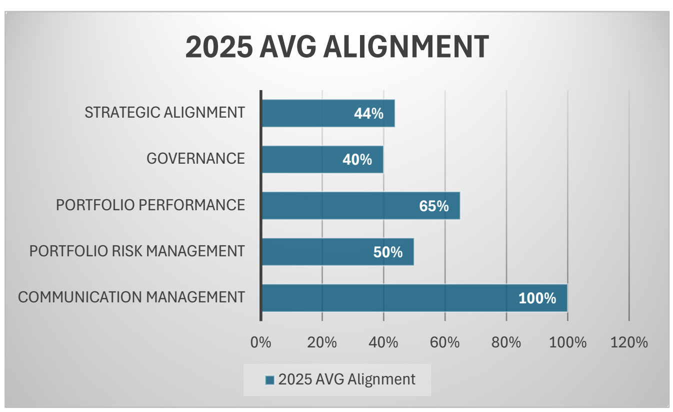

Figure 3. Average Alignment Scores

Table 4 depicts the detailed view of the analysis in the Strategic Alignment domain. Partial alignment existed in such tasks as evaluating organizational strategic goals, gathering data, and identifying potential portfolio components through business plans because those tasks must be done even in a program-centric model. There was no discernable alignment for two of the eight tasks because they spoke specifically to tasks carried out by an organization with the structure and policy to execute portfolio management.

Table 4. Strategic Alignment Domain Comparison

Tasks	Strategic Alignment (25%)	DOD UOC	DOD Competency	2025 Alignment	2025 BTI
Task 1	Evaluate organizational strategic goals and objectives using document reviews, interviewing, and other information gathering techniques in order to understand the strategic priorities.	AM1, L2	Acquisition Program Strategic Planning (AM1); Vision (L2)	0.5	1
Task 2	Identify prioritization criteria (e.g., legislative, dependencies, ROI, stakeholder expectations, strategic fit) using information gathering and analysis techniques in order to create a basis for decision making.	AM2, AM4, TM1	Acquisition Policy and Best Practices (AM2); Political Savvy and External Situational Awareness (AM4); Decision Analysis (TM1)	0.5	1
Task 3	Rank strategic priorities working with key stakeholders and using qualitative and quantitative analyses in order to provide a guiding framework to operationalize the organizational strategic goals and objectives.			0	1
Task 4	Identify existing and potential portfolio components by reviewing documentation such as business plans/proposals in order to create portfolio scenarios.			0	1
Task 5	Create portfolio scenarios (what-if analysis) by reviewing components against prioritization criteria and using analysis techniques (e.g., options analysis, risk analysis, SWOT analysis, financial analysis) in order to evaluate and select viable options.	AM3, L4	Risk/Opportunity Management (AM3); Program Oversight (AM3); Decisiveness (L4)	0.5	1
Task 6	Recommend portfolio scenario(s) and related components, based on prioritization analysis/criteria, in order to provide governance with a rationale for decision making.	L4	Decisiveness (L4)	0.5	1
Task 7	Determine the impact to portfolio and portfolio components due to changes in strategic goals and objectives, in order to sustain strategic alignment.	AM1,BM2	Acquisition Program Strategic Planning (AM1); Financial Planning (BM2)	1	0
Task 8	Create high level portfolio roadmap working with key stakeholders using prioritization, interdependency analysis, and organizational constraints in order to confirm and communicate the portfolio components sequencing, dependencies, and strategic alignment.	AM3, AM4, L1, L2	Resource Management (AM3); Political Savvy and External Situational Awareness (AM4); Communicate Effectively (L1); Vision (L2); Flexibility (L2)	0.5	2
			Average Score	43.75%	1.00

The most significant gaps in the DoD competency standard regarding portfolio management are related to the Governance domain with a 40% alignment (as shown in Table 5). The tasks in this domain include establishing policies, procedures, authorities, and management models that align with portfolio management practices. Within the current DoD standards, these governance models either do not exist or, at the very least, are not codified in writing.

Table 5. Governance Domain Comparison

Tasks	Governance (20%)	DOD UOC	DOD Competency	2025 Alignment	2025 BTI
Task 1	Define and establish a governance model including the structure (including but not limited to steering committees, governance boards), policies, and decision-making roles, responsibilities, rights and authorities in order to support effective decision-making and achieve strategic goals.			0	3
Task 2	Determine portfolio management standards, protocols, rules, and best practices, using organizational assets (such as information systems, subject matter experts) and industry standards in order to establish consistent portfolio management practices.	AM3, AM5	Program Oversight (AM3); Tailoring Acquisition Approach (AM5)	0.5	2
Task 3	Define and/or modify portfolio processes and procedures including but not limited to benefits realization planning, information management, performance, communication, risk management, stakeholder engagement, resource management, and change portfolio efficiently and effectively.management in order to manage the	AM5	Tailoring Acquisition Approach (AM5)	0.5	2
Task 4	Create the portfolio management plan including, but not limited to, roles and responsibilities, governance model, escalation procedures, risk tolerances, and governance thresholds, change control and management, key performance indicators, prioritization model, and communication procedures using standards, models, and other organizational assets in order to ensure effective and efficient portfolio management.	AM3	Program Oversight (AM3)	0.5	2
Task 5	Make recommendations and obtain approval regarding portfolio decisions (e.g, components, plans, budget, roadmap) through communication with key decision makers as defined by the governance model, in order to authorize the execution of the portfolio.	AM4, L1	Politial Savvy and External Situational Awareness (AM4); Pathway Selection and Tailoring Acquisition Approach (AM5); Communicate Effectively (L1)	0.5	2
			Average Score	40.00%	2.20

In the domain of Portfolio Performance, the DoD competency standard was 65% aligned with the PfMP standard. Full alignment was observed in four of the 10 tasks and partial alignment in five. As shown in Table 6, the places where the standards aligned include monitoring performance and ensuring strategic alignment with organizational goals. Moreover, they aligned in training personnel to escalate issues to appropriate decision-makers, propose solutions, and determine the decision's impacts on the organization. However, the standards did not align in one of the 10 tasks related to Portfolio Performance. Specifically, the PfMP standard calls for documenting portfolio artifacts. Since the DoD only trains personnel at the program level, portfolio-level documentation of approvals, prioritizations, and decisions remains a gap.

Table 6. Portfolio Performance Comparison

Tasks	Portfolio Performance (25%)	DOD UOC	DOD Competency	2025 Alignment	2025 BTI
Task 1	Initiate the portfolio using the portfolio roadmap and supporting artifacts in order to authorize the portfolio structure and activate the components.	AM1, AM5	Acquisition Program Strategic Planning (AM1); Pathway Selection and Tailoring Acquisition Approach (AM5)	0.5	2
Task 2	Collect and consolidate key performance metric data, as defined by portfolio governance and using various techniques, in order to measure the health of the portfolio.	AM3	Program Oversight (AM3)	0.5	1
Task 3	Monitor the portfolio performance on an ongoing basis, using reports, conversations, dashboards, and auditing techniques in order to ensure portfolio effectiveness and efficiency and maintain strategic alignment.	AM3	Program Oversight (AM3)	1	0
Task 4	Manage and escalate issues by communicating recommended actions to appropriate decision makers for timely approval and implementation of proposed solution(s).	L1	Communicate Effectively (L1)	1	0
Task 5	Manage portfolio changes using change management techniques, in order to improve portfolio performance and maintain strategic alignment.	AM1, AM5, L2, L4	Requirements Management (AM1); Acquisition Program Strategic Planning (AM1); Tailoring Acquisition Approach (AM5); Flexibity (L2); Problem Solving (L4)	1	0
Task 6	Balance portfolio and prioritize portfolio components, using established criteria and methods in order to optimize resource utilization and achieve strategic portfolio objectives.	AM1, AM3, L2	Acquistion Program Strategic Planning (AM1); Program Oversight (AM3); Flexibility (L2)	0.5	2
Task 7	Analyze and optimize the consolidated allocation/reallocation of capacity (e.g., people, tools, materials, technology, facilities, financial) using supply/demand management and scenario analysis techniques to ensure portfolio efficiency and effectiveness.	AM3	Requirements Decomposition (AM3); Program Oversight (AM3)	1	0
Task 8	Update and refine existing portfolio road maps, using change analysis in order to facilitate re-allocation of organizational resources to the portfolio.	AM3, L2	Requirements Decomposition (AM3); Program Oversight (AM3); Flexibility (L3)	0.5	2
Task 9	Measure the aggregated portfolio performance results against the defined business or strategic goals and objectives in order to demonstrate progress toward the achievement of business or strategic goals.	BM2	Financial Planning (BM2); Programming (BM2)	0.5	2
Task 10	Maintain records by capturing portfolio artifacts, such as approvals, prioritizations, and other decisions, in order to ensure compliance with organizational policies, regulatory requirements, and portfolio management standards.			0	2
			Average Score	65.00%	1.10

As depicted in Table 7, the Portfolio Risk Management domain was 50% aligned. The DoD standard devotes significant time to outlining ways in which acquisitions personnel must identify and mitigate risk. However, in half of the tasks listed in the PfMP standard, the document speaks directly to processes and procedures unique to a portfolio management structure. These include tasks such as dependency analysis, portfolio-level risk registers, and analysis of portfolio management reserves. The DoD's program-centric training does not require similar practices.

Table 7. Portfolio Risk Management Comparison

Tasks	Portfolio Risk Management (15%)	DOD UOC	DOD Competency	2025 Alignment	2025 BTI
Task 1	Determine acceptable level of risk for the portfolio, based on organizational and stakeholder risk tolerances, in order to provide input to governance.	AM3	Risk/Opportunity Management (AM3)	1	0
Task 2	Develop the portfolio risk management plan, using governance risk guidelines, processes, and procedures and other organizational assets in order to capitalize on opportunities, and respond to risks.	AM3	Risk/Opportunity Management (AM3)	1	0
Task 3	Perform dependency analysis to identify and monitor risks related to the interdependencies and intradependencies within or across portfolios in order to support decision-making.			0	1
Task 4	Develop, monitor, and maintain portfolio-level risk register, including risks to strategic goals and objectives, to business value, and escalated from portfolio components, using risk management processes in order to support decision making.	AM3	Risk/Opportunity Management (AM3)	1	0
Task 5	Promote common understanding and stakeholder ownership of portfolio risks, through communications with stakeholders, in order to facilitate risk response.			0	1
Task 6	Provide recommendation and obtain approval for a portfolio management reserve, based on aggregate portfolio risk exposure, in order to optimize portfolio strategic goals and objectives.			0	1
			Average Score	50.00%	0.50

Table 8 shows the alignment of the two standards in the domain of Communications Management with 100% alignment observed. The DoD standard goes to great lengths to describe the type of communication the DoD expects from its acquisition professionals. This training is easily transferrable to a portfolio management format. Moreover, in this section of the PfMP standard, there is less portfolio-specific verbiage used. Instead, the focus is on how portfolio managers engage stakeholders and communicate up and down the chain of command.

Table 8. Communications Domain Comparison

Tasks	Communications Management (15%)	DOD UOC	DOD Competency	2025 Alignment	2025 BTI
Task 1	Analyze internal and external stakeholders using techniques such as meetings, interviews, surveys/questionnaires, in order to identify stakeholder expectations, interests, and influence on the success of the portfolio.	AM3, AM4; L1	Program Oversight (AM3); Teaming (AM3); Political Savvy and External Situational Awareness (AM4); Communicate Effectively (L1)	1	0
Task 2	Create the aggregate communication strategy and plan, including methods, recipients, vehicles, timelines and frequencies in order to enable effective communication to stakeholders.	L1	Communicate Effectively (L1)	1	0
Task 3	Engage stakeholders, through oral and written communication, to ensure awareness, manage expectations, foster support, and build relationships and collaboration for the success of the portfolio roadmap.	AM3, AM4, L1	Program Oversight (AM3); Teaming (AM3); Political Savvy and External Situational Awareness (AM4); Communicate Effectively (L1)	1	0
Task 4	Maintain the communication strategy and plan by evaluating current communications capabilities, identifying gaps, and documenting communications plan to meet stakeholder requirements.	AM4; L1	Political Savvy and External Situational Awareness (AM4); Communicate Effectively (L1)	1	0
Task 5	Prepare and/or facilitate stakeholder understanding of portfolio management related processes, procedures, and protocols using organizational assets (e.g., information systems, training delivery methods) in order to promote common understanding and application of the portfolio management process.	AM4; L1	Political Savvy and External Situational Awareness (AM4); Communicate Effectively (L1)	1	0
Task 6	Verify accuracy, consistency, and completeness of portfolio communication, utilizing governance guidelines, to maintain credibility and satisfaction with all stakeholders.	AM4; L1	Political Savvy and External Situational Awareness (AM4); Communicate Effectively (L1)	1	0
			Average Score	100%	0.00

Figure 4 reflects the BTI rating for each domain of the PfMP standard. The average overall BTI score is 1, reflecting a low to medium BTI level for most gaps observed in the DoD standard. This means that many of the skills trained in the DoD PM standards are transferable to the PfM model with few modifications. However, one area where the transition will be difficult is in the domain of governance, where the researchers assessed a BTI rating of 2.2. Currently, DoD personnel structures, policies, and procedures are set for a program-centric model of governance. The DoD will need to modify personnel structure, current governance policies, and associated procedures toward a PfM-centric structure to transition to a PfM structure. Changes in the domain of governance will allow for changes across all domains analyzed in this research.

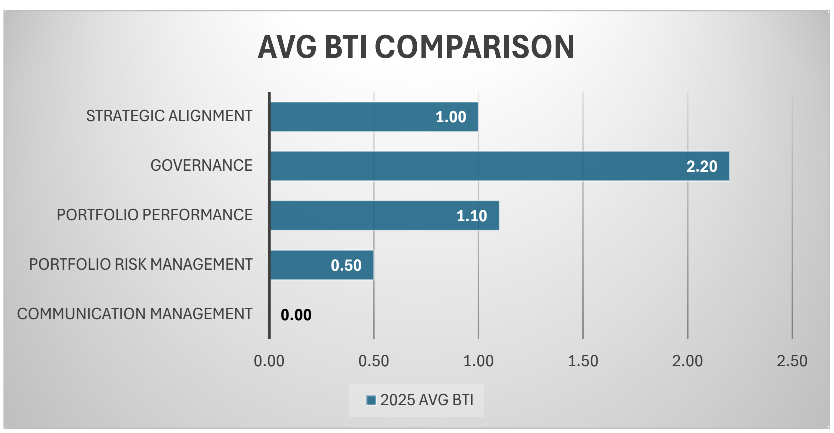

Figure 4. BTI Breakdown by PfMP Domain

Summary

The analysis indicates significant gaps in the DoD PM competency standards that must be addressed before the DoD can fully implement PfM. These findings are consistent with the recommendations from the Section 809 panel and GAO reports. Currently, DoD acquisition operates on a program-centric model that stovepipes funding into specific programs. Moreover, DoD PMs have little insight and influence into the acquisition program baselines of adjacent PMs within the same PEO or other PEOs (Shultz, 2020).

In the governance domain, the PfMP standard calls for personnel to “define and establish a governance model, policies, and decision-making roles” (PMI, 2013, p. 5). For the

DoD, this would require significant restructuring and policy reform. Most importantly, portfolio managers' authorities, roles, and responsibilities must be codified to incorporate the tasks outlined in the governance domain. Once the structure is in place, the PfMP standard outlines the need for each portfolio manager to enact a "portfolio management plan" (PMI, 2013, p. 5). This includes authoritative thresholds, risk tolerance levels, key performance indicators, prioritization models, and escalation procedures within each portfolio. While similar considerations exist inside many programs, the infrastructure does not currently exist at the portfolio level within the DoD.

The second domain in which the DoD has significant gaps in project management standards is strategic alignment. This PfMP domain calls for leaders to make and evaluate organizational goals and marry them to portfolios (PMI, 2013). Once the goals align with portfolios, the PfMP standard calls for portfolio managers to set prioritization criteria using analytical decision-making tools, resulting in a portfolio road map used to budget, plan, and execute. The PfMP standard calls for impact analysis of shortfalls within the portfolio road map (PMI, 2013).

Within the portfolio risk management domain, the current DoD competency standards capture the understanding, planning, and mitigating of risk thoroughly. However, adding the higher lens from the portfolio level is essential for effective portfolio risk management. In this regard, the DoD needs to continue to develop standards that capture this increased awareness of risk and how changes in one program can increase or decrease risks in an adjacent program within a portfolio. Under the current model, stove-piped programs often lack the proper coordination and awareness of adjacent programs.

Within the portfolio performance domain, some alignment was observed, specifically in tasks dealing with accountability, maintaining high standards, and making well-informed and timely decisions. These competencies are central to basic military standards and culture and are currently trained to and evaluated in PM competency standards. These tasks will carry over well to the PfM construct in the future.

The DoD and PMI standards were fully aligned in the domain of communications management. The tasks in this domain center around leadership, developing leaders, and developing rapport with vendors. Communications management competency is the strength that can enable forward momentum for the DoD to overcome BTIs to make swift and efficient progress towards transition.

Conclusions

The research and results show two significant conclusions:

- The DoD has yet to embrace portfolio management and meet the requirements laid out in FY2021 NDAA and several recommendations from the GAO and the Section 809 Panel.
- The DoD is more than capable of implementing portfolio management.

The research results suggest that the most significant BTIs reside in the governance domain. This is a result of the current program-centric construct called for by the Goldwater–Nichols Act that resulted in the basic governance construct still in place (Section 809 Panel, 2019a). It divides the acquisition governance into three decision support systems: requirements (formerly referred to as the Joint Capabilities Integration and Development System [JCIDS] for formal programs of record]; resourcing (Planning, Programming, Budgeting, and Execution [PPBE] system); and the Adaptive Acquisition Framework. Each of these decision support systems is fundamentally driven by different and often contradictory goals:

- The requirements generation system is driven primarily by a combination of capability needs and an evolving threat.
- The resource allocation system is calendar-driven, with an annual appropriations bill providing funding for acquisition efforts.
- The Adaptive Acquisition Framework is event-driven by milestones; it is based on commercial industry best practices of knowledge points and off-ramps supported by the design, development, and testing of the systems as technology, system design, and manufacturing processes mature.

The disjointed nature of this construct will be the most significant barrier to implementation of PfM. These findings are consistent with the Section 809 Panel's (2019a) analysis.

This analysis does not indicate that the DoD is incapable of conducting PfM. Instead, in conducting PfM, the DoD relies on PM competency standards that do not align with industry best practices. Defense acquisition is not currently structured to provide the appropriate training, evaluation, and feedback for proper job performance within a PfM-centric strategy. The establishment of PfM competencies remains a vital component to the successful implementation of congressional mandates to move toward a PfM-centric acquisition strategy.

The DoD should consider modifying its governance structure to recognize "portfolio manager" as an official career field. This is consistent with the Section 809 Panel recommendations, which assigned these responsibilities and authorities to portfolio acquisition executives (PAE; Section 809 Panel, 2019a). The PAE construct is analogous to the current PEO, except with expanded responsibilities and authorities. Concurrently, the Services should support acquisitions professionals obtaining PfMP certifications and include PfMP certification in the requirements for key acquisition positions.

Lastly, future research should address funding transfer authorities within defense acquisition and the establishment of portfolio elements for budgeting rather than program elements. Portfolio managers should be given milestone decision authority of assigned programs and projects and be allowed to manage cost, schedule, and performance within a portfolio acquisition baseline as opposed to acquisition program baselines.

This research indicates that the DoD may be able to implement changes necessary to implement portfolio management. This is supported by the BTI analysis of this study. Additionally, the Warfighting Acquisition University (WAU) course offerings are tailorable, and the PM competencies could be adjusted to align to PMI PfM standards. The DoD should consider adopting the PMI standards and certification processes as the baseline standards for the DoD, then make addendums or amendments to accommodate any differences between the DoD and industry when implementing changes.

Portfolio management requires a higher level of training and experience. For leaders to perform and be evaluated on PfM key domains properly, they must receive adequate training supported by clearly defined career-field competency models. Establishing PfMP competency standards will not fully resolve these shortfalls due to the various other policies and structural changes that require reform. However, training and evaluating acquisitions professionals on incorporating the proper aspects of PfMP competency domains will be essential to moving forward with a portfolio-centric approach.

The results of this research are aligned with recent Acquisition reform initiatives outlined the November 7, 2025, memorandum by the secretary of war titled, "Transforming the Defense Acquisition System into the Warfighting Acquisition System to Accelerate Fielding of Urgently Needed Capabilities to Our Warriors" (Office of the Secretary of War [OSW], 2025). In line with previous Section 809 recommendations, the memorandum directs the under secretary of war for acquisition and sustainment (USW(A&S)) to establish PAEs (OSW, 2025). The Military Services, through their service acquisition executives, will establish PAEs as the single accountable official for portfolio outcomes with authority to do the following:

1. Structure programs as schedule-driven capability increments with aggressive production delivery schedules, unit-cost ceiling goals, and broad mission effectiveness goals. Make trade-offs throughout the development to permit iterative enhancement and rapid delivery of subsequent increments. Make prudent cost, schedule, and performance trades that prioritize time-to-field, including execution of portfolio-level programming within defined and authorized boundaries.
2. Implement capability trade councils, replacing configuration steering boards, to integrate operational and acquisition authorities and make requirement trade-offs, and waive technical standards and other certification requirements not mandated by statute or safety.
3. Maximize use of Modular Open System Architectures (MOSA) for development programs moving forward by obtaining delivery of critical system interfaces with government purpose rights enable modular competition and supply chain resiliency.
4. Organizationally align contracting officers to report directly to PAEs in the acquisition chain of command, to maintain responsiveness to the operational problems that materiel solutions are intended to address (Secretary of War, 2025).

References

Ahern, D., & Driessnack, J. (2019, June 30). *Implementing a multitier portfolio management structure for defense acquisition.* Olde Stone Consulting. http://www.oldestoneconsulting.com/blog/2019/6/30/implementing-a-multitier-portfolio-management-structure-for-defense-acquisition

American National Standards Institute. (n.d.-a). *ANSI introduction*. Retrieved May 7, 2025, from https://www.ansi.org/about/introduction

American National Standards Institute. (n.d.-b). *Standards FAQs.* Retrieved May 7, 2025, from https://www.ansi.org/standards-faqs#:~:text=A%20standard%20is%20a%20document,best%20way%20of%20doing%20something

American Society for Quality. (n.d.). *Standards 101|ASQ*. Retrieved May 7, 2025, from https://asq.org/quality-resources/standards-101#standards

Anton, P., McKernan, M., Munson, K., Kallimani, J., Levedahl, A., Blickstein, I., Drezner, J., & Newberry, S. (2019). *Assessing Department of Defense use of data analytics and enabling data management to improve acquisition outcomes* (RR-3136-OSD). RAND Corporation. https://www.rand.org/pubs/research_reports/RR3136.html

Biedenbach, T., & Müller, R. (2012). Absorptive, innovative, and adaptive capabilities and their impact on project and project portfolio performance. *International Journal of Project Management*, *30*(5), 621–635. https://doi.org/10.1016/j.ijproman.2012.01.016

Distel, D., Hannon, E., Krause, M., & Krieg, A. (2020, December 4). *Finding the sweet spot in product portfolio management.* McKinsey & Company. https://www.mckinsey.com/business-functions/operations/our-insights/finding-the-sweet-spot-in-product-portfolio-management

Driessnack, J. D., & Johnson, J. (2023). *Portfolio performance analysis and visualization* (AIRC-2023-TR-012). Acquisition Innovation Research Center. https://apps.dtic.mil/sti/trecms/pdf/AD1214875.pdf

Ewing, L., Dell, R., MacCalman, M., & Whitney, L. (2013, January). *Capability portfolio analysis tool (CPAT) verification and validation report* (NPS-OR-13-001). Naval Postgraduate School. https://calhoun.nps.edu/handle/10945/27340

Gates, S. M., Esteves, F., Roth, E., & Kempf, J. (2024). *Implementation of the new Defense Acquisition Workforce Improvement Act framework: End of fiscal year 2022 update* (RRA768-3). RAND. https://www.rand.org/pubs/research_reports/RRA758-3.html

GAO. (2020, June 3). *Defense acquisitions annual assessment: Drive to deliver capabilities faster increases importance of program knowledge and consistent data for oversight* (GAO-20-439). https://www.gao.gov/products/gao-20-439

Gutiérrez, E., & Magnusson, M. (2014). Dealing with legitimacy: A key challenge for project portfolio management decision makers. *International Journal of Project Management*, *32*(1), 30–39. https://doi.org/10.1016/j.ijproman.2013.01.002

Heising, W. (2012). The integration of ideation and project portfolio management: A key factor for sustainable success. *International Journal of Project Management*, *30*(5), 582–595. https://doi.org/10.1016/j.ijproman.2012.01.014

International Standards Organization. (n.d.-a). *ISO—About ISO*. Retrieved May 7, 2025, from https://www.iso.org/about

International Standards Organization. (n.d.-b). *ISO 9000:2015(en), Quality management systems—Fundamentals and vocabulary.* Retrieved May 7, 2025, from https://www.iso.org/obp/ui/#iso:std:iso:9000:ed-4:v1:en

Janiga, M., & Modigliani, P. (2014, November–December). Think portfolios, not programs. *Defense AT&L Magazine*, *43*(6), 12–16. https://dair.nps.edu/bitstream/123456789/3013/1/SEC809-RSC-14-0038.pdf

Karnes, J. L. (2020). *Aligning DoD program management competencies with the project management institute standards* [Master's thesis, Naval Postgraduate School]. Defense Acquisition Innovation Repository: https://dair.nps.edu/handle/123456789/4298

Kock, A., Schulz, B., Kopmann, J., & Gemünden, H. G. (2020). Project portfolio management information systems' positive influence on performance: The importance of process maturity. *International Journal of Project Management*, *38*(4), 229–241. https://doi.org/10.1016/j.ijproman.2020.05.001

Li, Y. M., Chen, H. M., Liou, J. H., & Lin, L. F. (2014). Creating social intelligence for product portfolio design. *Decision Support Systems*, *66*, 123–134. https://doi.org/10.1016/j.dss.2014.06.013

Mak, M. A. (2022). *Tactical aircraft investments: DoD needs additional portfolio analysis to inform future budget decisions* (GAO-23-106375). United States Government Accountability Office. https://www.gao.gov/products/gao-23-106375

Modigliani, P. (2015, April 1). *Portfolio acquisition: How DOD can leverage the commercial product line model.* Naval Postgraduate School. https://calhoun.nps.edu/handle/10945/53650

National Defense Authorization Act for Fiscal Year 2020, Pub. L. No. 116-92, Sec. 861 (2019). https://www.congress.gov/116/bills/s1790/BILLS-116s1790enr.pdf

National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116–283, Sec. 836. (2021). https://www.congress.gov/116/bills/hr6395/BILLS-116hr6395enr.pdf

Nippa, M., Pidun, U., & Rubner, H. (2011). Corporate portfolio management: Appraising four decades of academic research. *Academy of Management Perspectives*, *25*(4), 50–66. https://doi.org/10.5465/amp.2010.0164

Oakley, S. S. (2021). *Weapon systems annual assessment: Updated program oversight approach needed* (GAO-21-222). United States Government Accountability Office. https://www.gao.gov/products/gao-21-222

Oakley, S. S. (2025). *Heightened attention could save billions more and improve government efficiency and effectiveness* (GAO-25-1077430). United States Government Accountability Office. https://www.gao.gov/products/gao-25-107743

Office of the Assistant Secretary of Defense. (2024, December 16). *Program management functional career field competencies* [Memorandum]. Assistant Secretary of Defense for Acquisition.

Office of the Under Secretary of Defense for Acquisition and Sustainment. (2022, March 25). *Defense acquisition workforce education, training, experience, and career development program* (DoD Instruction 5000.66, Change 3). Department of Defense. https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500066p.PDF

Office of the Under Secretary of Defense for Acquisition and Sustainment. (2023, September 25). *Capability portfolio management* (DoD Directive 7045.20). Department of Defense. https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/704520p.pdf

Office of the 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]. Department of War. 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 U.S. Congressman Joe Wilson. (2020, July 7). *DOD Portfolio Management Accountability Act included in National Defense Authorization Act* [Press release]. https://joewilson.house.gov/media-center/press-releases/dod-portfolio-management-accountability-act-included-in-ndaa

Petit, Y. (2012). Project portfolios in dynamic environments: Organizing for uncertainty. *International Journal of Project Management*, *30*(5), 539–553. https://doi.org/10.1016/j.ijproman.2011.11.007

Project Management Institute. (n.d.). *What are project management standards?* Retrieved May 7, 2025, from https://www.pmi.org/standards/about

Project Management Institute. (2013). *Portfolio Management Professional (PfMP) examination content outline.*

Project Management Institute. (2017a). *Portfolio Management Professional (PfMP) handbook.*

Project Management Institute. (2017b). *The standard for portfolio management* (4th ed.).

PMO Advisory. (n.d.). *Portfolio Management PfMP–PMO Advisory*. Retrieved May 7, 2025, from https://www.pmoadvisory.com/pmi-certification/portfolio-management-pfmp/

Section 809 Panel. (2017, May). *Advisory panel on streamlining and codifying acquisition regulations: Section 809 Panel interim report.* Defense Technical Information Center. https://discover.dtic.mil/wp-content/uploads/809-Panel-2019/Interim-Report/Sec809Panel_Interim-Report_May2017.pdf

Section 809 Panel. (2018a, January). *Report of the advisory panel on streamlining and codifying acquisition regulations: Volume 1 of 3.* Defense Technical Information Center. https://discover.dtic.mil/wp-content/uploads/809-Panel-2019/Volume1/Sec809Panel_Vol1-Report_Jan2018.pdf

Section 809 Panel. (2018b, June). *Report of the advisory panel on streamlining and codifying acquisition regulations: Volume 2 of 3.* Defense Technical Information Center. https://discover.dtic.mil/wp-content/uploads/809-Panel-2019/Volume2/Sec809Panel_Vol2-Report_Jun2018.pdf

Section 809 Panel. (2019a, January). *Report of the advisory panel on streamlining and codifying acquisition regulations: Volume 3 of 3 part 1.* Defense Technical Information Center. https://discover.dtic.mil/wp-content/uploads/809-Panel-2019/Volume3/Sec809Panel_Vol3-Report_Jan2019_part-1_0509.pdf

Section 809 Panel. (2019b, January). *Report of the advisory panel on streamlining and codifying acquisition regulations: Volume 3 of 3 part 2.* Defense Technical Information Center. https://discover.dtic.mil/wp-content/uploads/809-Panel-2019/Volume3/Sec809Panel_Vol3-Report_Jan2019_part-2_0307.pdf

Section 809 Panel. (2019c, February). *Advisory panel on streamlining and codifying acquisition regulations: A roadmap to the Section 809 Panel reports.* Federal Depository Library Program. https://permanent.fdlp.gov/websites/discover.dtic.mil/discover.dtic.mil/discover.dtic.mil/wp-content/uploads/809-Panel-2019/Roadmap/Sec809Panel_Roadmap_DEC2019.pdf

Shultz, B. (2020, March). A portfolio management–based acquisition model? *Contract Management*, *60*(3), 42–47. https://www-proquest-com.libproxy.nps.edu/docview/2458099153?accountid=12702

Stewart, C. W., Deitrich, A. T., & Reid, J. M. (2021). *Gap analysis of Department of Defense program management competency standards in preparation for the shift to portfolio management in defense acquisition* [Master's thesis, Naval Postgraduate School]. NPS Archive: Calhoun. https://dair.nps.edu/handle/123456789/4517

Sullivan, M. J. (2015). *Weapon system acquisitions: Opportunities exist to improve the Department of Defense's portfolio management* (GAO-15-466). United States Government Accountability Office. https://www.gao.gov/products/gao-15-466

Sullivan, M. J. (2018). *Defense acquisition workforce: Opportunities exist to improve practices for developing program managers* (GAO-18-217). United States Government Accountability Office. https://www.gao.gov/products/gao-18-217

Thompson, R., & Johnson, M. (2019, July 10). 809 Panel calls for managing "capabilities," not weapons. *Breaking Defense.* https://breakingdefense.com/ 2019/07/809-panel-calls-for-managing-capabilities-not-weapons/

Young, M., & Conboy, K. (2013). Contemporary project portfolio management: Reflections on the development of an Australian competency standard for project portfolio management. *International Journal of Project Management*, *31*(8), 1089–1100. https://doi.org/10.1016/j.ijproman.2013.03.005

Appendix

PMI Portfolio Domain Tasks

Domain 1: Strategic Alignment. The purpose of the Strategic Alignment domain is to evaluate an individual's ability to align all components that make up a portfolio, including programs and projects, to the organization's overall strategic objectives and priorities (PMI, 2013). This highlights portfolio management's focus on strategic management. The Strategic Alignment and Portfolio Performance domains are the most heavily weighted portions of the exam at 25% each. The Strategic Alignment domain contains eight tasks, as listed in Table 9.

Table 9. Domain 1: Strategic Alignment Tasks.
(PMI, 2013, p. 4)

Tasks	Strategic Alignment (25%)
Task 1	**Evaluate organizational strategic goals and objectives using document reviews, interviewing, and other information gathering techniques in order to understand the strategic priorities.**
Task 2	**Identify prioritization criteria (e.g., legislative, dependencies, ROI, stakeholder expectations, strategic fit) using information gathering and analysis techniques in order to create a basis for decision making.**
Task 3	**Rank strategic priorities working with key stakeholders and using qualitative and quantitative analyses in order to provide a guiding framework to operationalize the organizational strategic goals and objectives.**
Task 4	**Identify existing and potential portfolio components by reviewing documentation such as business plans/proposals in order to create portfolio scenarios.**
Task 5	**Create portfolio scenarios (what-if analysis) by reviewing components against prioritization criteria and using analysis techniques (e.g., options analysis, risk analysis, SWOT analysis, financial analysis) in order to evaluate and select viable options.**
Task 6	**Recommend portfolio scenario(s) and related components, based on prioritization analysis/criteria, in order to provide governance with a rationale for decision making.**
Task 7	**Determine the impact to portfolio and portfolio components due to changes in strategic goals and objectives, in order to sustain strategic alignment.**
Task 8	**Create high level portfolio roadmap working with key stakeholders using prioritization, interdependency analysis, and organizational constraints in order to confirm and communicate the portfolio components sequencing, dependencies, and strategic alignment.**

Domain 2: Governance. The purpose of the Governance domain is to evaluate an individual's ability to oversee the portfolio; to create the overall management plan, including performance standards, best practices, processes and procedures, and overall management structure; and to manage decision-making elements to ensure proper authorization of portfolio execution (PMI, 2013). The Governance domain, weighted at 20%, is the third most important set of competencies behind Strategic Alignment and Portfolio Performance. It includes the five tasks listed in Table 10.

Table 10. Domain 2: Governance Tasks.
(PMI, 2013, p. 5)

Tasks	Governance (20%)
Task 1	Define and establish a governance model including the structure (including but not limited to steering committees, governance boards), policies, and decision-making roles, responsibilities, rights and authorities in order to support effective decision-making and achieve strategic goals.
Task 2	Determine portfolio management standards, protocols, rules, and best practices, using organizational assets (such as information systems, subject-matter experts) and industry standards in order to establish consistent portfolio management practices.
Task 3	Define and/or modify portfolio processes and procedures including but not limited to benefits realization planning, information management, performance, communication, risk management, stakeholder engagement, resource management, and change management in order to manage the portfolio efficiently and effectively.
Task 4	Create the portfolio management plan including, but not limited to, roles and responsibilities, governance model, escalation procedures, risk tolerances, and governance thresholds, change control and management, key performance indicators, prioritization model, and communication procedures using standards, models, and other organizational assets in order to ensure effective and efficient portfolio management.
Task 5	Make recommendations and obtain approval regarding portfolio decisions (e.g, components, plans, budget, roadmap) through communication with key decision makers as defined by the governance model, in order to authorize the execution of the portfolio.

Domain 3: Portfolio Performance. The purpose of the Portfolio Performance domain is to evaluate an individual's ability to oversee the execution of the portfolio within the established governance parameters set under the previous domain, to assess and balance the components of the portfolio based on performance and changes in strategic alignment, and to monitor the overall health of the portfolio (PMI, 2013). The Portfolio Performance domain, along with Strategic Alignment, is weighted at 25%. It includes the 10 tasks listed in Table 11.

Table 11. Domain 3: Portfolio Performance Tasks.
(PMI, 2013, p. 6)

Tasks	Portfolio Performance (25%)
Task 1	**Initiate the portfolio using the portfolio roadmap and supporting artifacts in order to authorize the portfolio structure and activate the components.**
Task 2	**Collect and consolidate key performance metric data, as defined by portfolio governance and using various techniques, in order to measure the health of the portfolio.**
Task 3	**Monitor the portfolio performance on an ongoing basis, using reports, conversations, dashboards, and auditing techniques in order to ensure portfolio effectiveness and efficiency and maintain strategic alignment.**
Task 4	**Manage and escalate issues by communicating recommended actions to appropriate decision makers for timely approval and implementation of proposed solution(s).**
Task 5	**Manage portfolio changes using change management techniques, in order to improve portfolio performance and maintain strategic alignment.**
Task 6	**Balance portfolio and prioritize portfolio components, using established criteria and methods in order to optimize resource utilization and achieve strategic portfolio objectives.**
Task 7	**Analyze and optimize the consolidated allocation/reallocation of capacity (e.g., people, tools, materials, technology, facilities, financial) using supply/demand management and scenario analysis techniques to ensure portfolio efficiency and effectiveness.**
Task 8	**Update and refine existing portfolio road maps, using change analysis in order to facilitate re-allocation of organizational resources to the portfolio.**
Task 9	**Measure the aggregated portfolio performance results against the defined business or strategic goals and objectives in order to demonstrate progress toward the achievement of business or strategic goals.**
Task 10	**Maintain records by capturing portfolio artifacts, such as approvals, prioritizations, and other decisions, in order to ensure compliance with organizational policies, regulatory requirements, and portfolio management standards.**

Domain 4: Portfolio Risk Management. The purpose of the Portfolio Risk Management domain is to evaluate an individual's ability to evaluate portfolio risk and align it with the risk appetite of the organization (PMI, 2013). It is weighted at 15% and includes the six tasks listed in Table 12.

Table 12. Domain 4: Portfolio Risk Management Tasks.
(PMI, 2013, p. 7)

Tasks	Portfolio Risk Management (15%)
Task 1	Determine acceptable level of risk for the portfolio, based on organizational and stakeholder risk tolerances, in order to provide input to governance.
Task 2	Develop the portfolio risk management plan, using governance risk guidelines, processes, and procedures and other organizational assets in order to capitalize on opportunities, and respond to risks.
Task 3	Perform dependency analysis to identify and monitor risks related to the interdependencies and intradependencies within or across portfolios in order to support decision-making.
Task 4	Develop, monitor, and maintain portfolio-level risk register, including risks to strategic goals and objectives, to business value, and escalated from portfolio components, using risk management processes in order to support decision making.
Task 5	Promote common understanding and stakeholder ownership of portfolio risks, through communications with stakeholders, in order to facilitate risk response.
Task 6	Provide recommendation and obtain approval for a portfolio management reserve, based on aggregate portfolio risk exposure, in order to optimize portfolio strategic goals and objectives.

Domain 5: Communications Management. The purpose of the Communications Management domain is to evaluate an individual's ability to conduct activities including stakeholder management, conflict management, and stakeholder engagement (PMI, 2013). It is weighted at 15% and includes the six tasks listed in Table 13.

Table 13. Domain 5: Communications Management Tasks.
(PMI, 2013, p. 8)

Tasks	Communications Management (15%)
Task 1	**Analyze internal and external stakeholders using techniques such as meetings, interviews, surveys/questionnaires, in order to identify stakeholder expectations, interests, and influence on the success of the portfolio.**
Task 2	**Create the aggregate communication strategy and plan, including methods, recipients, vehicles, timelines and frequencies in order to enable effective communication to stakeholders.**
Task 3	**Engage stakeholders, through oral and written communication, to ensure awareness, manage expectations, foster support, and build relationships and collaboration for the success of the portfolio roadmap.**
Task 4	**Maintain the communication strategy and plan by evaluating current communications capabilities, identifying gaps, and documenting communications plan to meet stakeholder requirements.**
Task 5	**Prepare and/or facilitate stakeholder understanding of portfolio management-related processes, procedures, and protocols using organizational assets (e.g., information systems, training delivery methods) in order to promote common understanding and application of the portfolio management process.**
Task 6	**Verify accuracy, consistency, and completeness of portfolio communication, utilizing governance guidelines, to maintain credibility and satisfaction with all stakeholders.**

PANEL 20. SUSTAINING SUPERIORITY: REINVIGORATING READINESS AND GLOBAL MARITIME POWER

Thursday, May 7, 2026, 1315 – 1430 ET (1015 - 1130 PT)

Panel Summary:

Maintaining a credible deterrent requires more than just acquiring new platforms; it demands a radical rethink of how the Department of War (DOW) sustains its existing arsenal and leverages global industrial capacity. This panel addresses the persistent decline in weapon system readiness by identifying root causes—such as neglected maintenance planning and fragmented data rights—and proposing a framework for predictive, lifecycle-centric management. Complementing this is a strategic analysis of international shipbuilding cooperation, evaluating pathways such as allied investment in U.S. yards and modular co-production to revitalize the maritime industrial base. Together, these research efforts offer a roadmap for achieving operational overmatch through resilient sustainment and strengthened allied partnerships.

Chair: CAPT Cedric McNeal, USN, Executive Director for Amphibious, Auxiliary, and Sealift (Acting)

Panel Presenters:

Evaluating Pathways for U.S. Shipbuilding Cooperation With Allies – *Henry H. Carroll, Research Associate, Center for Strategic & International Studies (CSIS)*

Shipbuilding Procurement: An International Analysis of Source Selection Processes – *Dr. Rene G. Rendon, Professor, Naval Postgraduate School*

Neglect with the Old, in with the New: How Neglect of Existing Weapon Systems Is Causing a Readiness Issue Within the Department of War, And How Acquisition Approaches Can be Improved to Solve It – *Moshe Schwartz, President, Etherton & Associates*

CAPT Cedric McNeal, USN—Born and raised in Jackson, Mississippi, Captain McNeal received his commissioning in 1997 from the Southern University and A&M College in Baton Rouge, Louisiana Naval Reserve Officers Training Corps (NROTC) unit. After commissioning, he reported to Surface Warfare Officer School in Newport, Rhode Island where he completed the basic Surface Warfare Officer training course. In November 1997, he reported aboard USS Gunston Hall (LSD 44) where served as Communications Officer and CMS Custodian/Electronic Key Management System (EKMS) Manager, qualifying as a Surface Warfare Officer, and also earning his qualification as a Diesel Engineering Officer of the Watch (EOOW). In November 1999, Captain McNeal reported aboard USS Caron (DD 970) and served as the ship's Navigator and Administrative Officer. Prior to completing his tour onboard Caron, Captain McNeal decided to pursue his commissioning option for the Engineering Duty Officer Community and in 2001, reported to the Naval Postgraduate School in Monterey, California. While there, he completed his Postgraduate studies, earning a Master of Science degree in Applied Physics with a concentration in Weapons Systems. His postgraduate studies were followed by his completion of the Engineering Duty Officer Basic Course in Port Hueneme, California.

Captain McNeal's first Engineering Duty Officer tour was with Supervisor of Shipbuilding (SUPSHIP) San Diego (now Southwest Regional Maintenance Center – SWRMC), where he was assigned as Project Officer to the USS Curts (FFG 38) for their Docking Selected Restricted Availability and then went on to serve in the same capacity for USS Peleliu (LHA 5). Also, while at SWRMC, Captain McNeal served as Project Manager for maintenance availabilities aboard USS Germantown (LSD 42) and USS Pearl Harbor (LSD 52) and then as the LSD/LPD Class Team Business Officer for all San Diego based LSD and LPD ships. Upon completion of his tour at SWRMC, Captain McNeal went on to complete two tours in new ship construction on the Gulf Coast. From December 2005 to May 2009, he was assigned to SUPSHIP Gulf Coast as the Production Officer for USS Makin Island (LHD 8). In June 2009, he reported to SUPSHIP Bath and led its Mobile, Alabama detachment in the delivery of USS Independence (LCS 2), a first of class vessel being presented to the Navy by a first-time shipbuilder.

Evaluating Pathways for U.S. Shipbuilding Cooperation With Allies

Henry H. Carroll—is a research associate with the Center for the Industrial Base at the Center for Strategic and International Studies (CSIS). His analytical focuses include the defense industrial base, naval policy, and assessing foreign military industrial capacity. Prior to joining CSIS, he worked as an intern in the defense and international practices of Brownstein and as a defense-focused legislative intern for Senate Majority Leader Chuck Schumer. Henry holds a BA in history, with a concentration in war and society, from Yale University. His undergraduate senior thesis examined the politics of U.S. naval shipbuilding in the interwar period. [hcarroll@csis.org].

Cynthia R. Cook—is a senior fellow in the Defense and Security Department at the Center for Strategic and International Studies (CSIS). Dr. Cook is a member of the editorial board for the Defense Acquisition Research Journal and is an adjunct professor at the Pardee RAND Graduate School. From 1997 to 2021, Dr. Cook worked as a senior management scientist at RAND, where she served as the director of the Acquisition and Technology Policy Center. She holds a PhD in sociology from Harvard University and a BS in management from the Wharton School of the University of Pennsylvania. [ccook@csis.org]

Seamus P. Daniels—is a fellow for Defense Budget Analysis in the Defense and Security Department at the Center for Strategic and International Studies (CSIS), where he researches issues related to U.S. and global defense funding, force structure, and military readiness. Prior to joining CSIS, Mr. Daniels worked for Government Executive Media Group. He holds an MA in international relations from Johns Hopkins University School of Advanced International Studies and an AB from Princeton University's School of Public and International Affairs with minors in Near Eastern studies and Arabic language and culture. [sdaniels@csis.org]

Oliver B. Buntin— is a research assistant with the Center for the Industrial Base at the Center for Strategic and International Studies. His research portfolio includes shipbuilding cooperation, landpower doctrine, supply chain issues, and the industrial base. Oliver is currently pursuing degrees in political science and architecture at the University of Chicago. [obuntin@csis.org]

Abstract

The U.S. naval shipbuilding industrial base faces well-known challenges delivering ships on time and at cost. The U.S. shipbuilding challenge is a complex, enterprise-wide issue. There is no single policy solution or silver bullet to solve the overall problem. Solutions are urgently needed, but the inherent policy trade-offs must be carefully weighed.

One option to address the challenges is through partnerships with close allies to enhance the shipbuilding enterprise. This research assesses three pathways for cooperation, including allied purchase and revitalization of U.S. shipyards; various methods of coproduction, including modular construction; and U.S. purchase of allied-built ships. These pathways, identified in a previous work, each present their own opportunities and challenges that are analyzed in this report. The paper also explored hybrid approaches that involve multiple pathways. The research focuses on the strong shipbuilding nations—and U.S. allies—South Korea and Japan as the most likely partners, but the findings are relevant to shipbuilding cooperation with other nations as well. This paper is a preliminary excerpt from a forthcoming work that discusses four pathways (including maintenance cooperation) and hybrid approaches in more detail.

Introduction

Once the world's preeminent shipbuilder, the United States now struggles to produce naval ships at the scale, speed, and cost its security requirements demand. During and after World War II, American shipyards delivered warships and commercial vessels in numbers that contributed to both its economic strength and maritime dominance. During the war, the United States produced nearly double the number of ships as the next five producers combined (Silva, 2024). That capacity has eroded over decades, leaving a narrower and more fragile industrial base. Today, workforce shortages, aging infrastructure, and brittle supply chains coupled with an uneven demand signal have weakened the naval shipbuilding enterprise (Daniels et al., 2025). Ships are routinely delivered years late and billions of dollars over initial estimates, constraining fleet readiness and consuming resources intended for modernization and force expansion.

The decline of U.S. shipbuilding reflects structural challenges that have accumulated over decades, limiting the effectiveness of purely domestic or unilateral solutions. An earlier CSIS assessment concluded,

> Despite the Navy's plans for growing the fleet and bipartisan efforts and funding from Congress, the U.S. shipbuilding enterprise—including the Navy, Department of Defense (DoD), Congress, and industry—has failed to consistently produce ships at the scale, speed, and cost demanded. These longstanding challenges stem from a series of interwoven, systemic issues within both the U.S. government and industry, as well as broader socioeconomic trends. (Daniels et al., 2025)

Unstable demand signals, evolving requirements, acquisition practices that diffuse accountability, limited competition in key segments, aging infrastructure, brittle supply chains, and a shrinking skilled workforce all contribute to persistent underperformance.

This report examines one policy solution that has gained increasing attention: cooperation with trusted allies to enhance the U.S. naval shipbuilding industrial base. It focuses on opportunities for deeper industrial integration with South Korea and Japan, two treaty allies that possess globally competitive commercial and naval shipbuilding sectors, advanced production techniques, and supplier ecosystems. Both countries operate high-throughput shipyards and have demonstrated the ability to deliver complex vessels on predictable schedules.

This report presents three pathways for cooperation, including allied purchase of, partnership with, or investment in yards in the United States; coproduction with allied shipbuilders; and the purchase of ships produced in allied yards overseas. It also addresses hybrid cases where multiple pathways are incorporated. The analysis assesses how these different pathways related to metrics including cost, speed, and development of indigenous shipbuilding capacity, and reviews their viability from the perspective of allied willingness and interest along with the U.S. regulatory and policy framework.

This report is not merely a theoretical exercise. Given global tensions and a rapidly rising U.S. shipbuilding budget, the United States and the Trump administration want to build ships quickly. At the same time, the governments of South Korea and Japan have excellent shipbuilding capabilities and seek to bolster cooperation with the United States to support their domestic industries and strengthen bonds with their most important military ally. This report lays out the ways U.S. policymakers can advance their goals via cooperation.

Pathways for Enhancing U.S. Shipbuilding Capability and Capacity

Carefully designed cooperation among U.S. industry, government, and trusted allies offers a practical approach to rebuild capacity while reinforcing collective maritime security. This report outlines three potential pathways for advancing allied shipbuilding integration, expanding on previous work on the topic (Carroll & Cook, 2025). Together, these pathways offer distinct mechanisms for addressing the U.S. naval shipbuilding challenge.

The three pathways explored in this paper are:

- Allied acquisition of, partnership with, or investment in U.S. shipyards to expand production capacity.

- Coproduction of warships by U.S. and allied shipyards, leveraging complementary industrial strengths.
- U.S. purchase of warships from allied shipyards to accelerate fleet expansion and reduce schedule risk.

In practice, these pathways are not mutually exclusive and have been frequently operationalized into a "hybrid cooperation approach," which is also discussed below. Moreover, other pathways also exist, such as expanding maintenance, repair, and overhaul (MRO) activities in allied shipyards overseas with the goal of repurposing U.S. maintenance facilities into new build yards. Collaboration on MRO will be considered in full in a forthcoming paper.

Criteria for Evaluating Cooperation Pathways

This report assesses these three different pathways by a series of metrics and reviews their viability from the perspective of allied willingness to cooperate and existing U.S. regulatory and policy framework. Each pathway is assessed against five principal criteria, grouped into those related to their impact on outcomes in impacting U.S. policy objectives and those related to the potential implementation of each pathway.

Outcome-related criteria include:

- *Ship delivery speed* assesses how a pathway may deliver ships expediently to the U.S. Navy.
- *Cost implications* assesses the budgetary repercussions of a given pathway.
- *Bolstering U.S. shipbuilding capacity* evaluates the extent to which a given pathway may expand the United States' indigenous shipbuilding capacity.

Implementation criteria include:

- *Allied industry willingness and ability* measures the extent to which industry partners in allied countries would be able and willing to cooperate with the United States in this manner.
- *U.S. political and regulatory viability* evaluates the palatability of each pathway within the U.S. shipbuilding enterprise given the interests of the Navy, broader DOD, Congress, and the shipbuilding industrial base as well as the requisite legal and regulatory changes needed for implementation.

The report identifies examples of ongoing or past cooperation. Some feature aspects that fall under multiple pathways. The final section of the report discusses potential models for future, hybrid approaches to cooperation.

The evaluation is not intended to select an optimal pathway for policymakers and the U.S. shipbuilding enterprise to implement. Rather, the assessment highlights which pathways correspond to U.S. policymaker preferences, such as the ability of the United States to rapidly expand the size of its fleet or to grow U.S. domestic shipbuilding capacity.

The pathways' evaluations are based on publicly available information as well as information gathered by the study team in interviews with current and former government officials, industry representatives, and think tank experts from the United States, Republic of Korea, and Japan. Interviews were conducted both in person, including during a visit to Korea and Japan in June and July 2025, and virtually.

Allied Purchase of, Partnership With, or Investment in U.S. Shipyards

The first pathway for cooperation involves allied builders either acquiring or partnering with U.S. shipyards to improve their construction capacity and efficiency. The distinguishing feature of this approach is direct investment in U.S. facilities. Direct acquisition occurs when a foreign shipbuilder purchases a U.S. yard outright, gaining full operational control and access to its expertise, supply chain relationships, and intellectual property. This access facilitates technology and management-practice transfer, by which tacit knowledge and proprietary processes, not just documented procedures, can be shared. It may also enable transformative facility investments from the allied shipbuilder. The principal constraints are regulatory and workforce-related: Acquisitions of yards involved in sensitive naval work typically require CFIUS review and other clearances or waivers, while the acquiring firm must reconcile its management practices with U.S. labor, subcontractor, and cultural norms. Foreign shipbuilding firms also need to acclimatize themselves to U.S. government contracting practices.

Foreign shipbuilders have already taken significant actions along this pathway. Italy's Fincantieri SpA. bought Manitowoc Marine Group in Wisconsin in 2009; Australian ferry-maker Austal bought out Bender Shipbuilder and Repair Company's stake in their joint venture in Alabama in 2006; and South Korea's Hanwha Ocean purchased Philly Shipyard in 2024.

Partnership is a broader category that encompasses joint ventures, minority investments, commercial design collaborations, technical assistance agreements, and personnel exchange programs. While the foreign entity does not take ownership of a U.S. yard in a partnership, knowledge moves through negotiated channels, such as shared designs, jointly developed technology, or workers rotating between facilities to develop hands-on skills. Compared to a commercial design license agreement, which transfers only what is explicitly documented and tends to be weaker for process-level learning, partnerships can approach the knowledge-transfer depth of acquisition if they include sustained on-site rotations and engagement.

Partnership arrangements generally face lower regulatory scrutiny than acquisitions but must comply with ITAR, Export Regulations, and DSCA guidelines. Among the most noteworthy of these arrangements, South Korea's Daewoo Marine Engineering and Hyundai Heavy Industries have initiated partnerships with General Dynamics NASSCO and Huntington Ingalls Industries, respectively.

Pathway Evaluation

Ship Delivery Time

The speed with which foreign acquisitions or partnerships may lead to a production uptick depends on the parameters of the chosen approach, demand signal, regulatory barriers, and yard-specific physical and workforce conditions.

Direct acquisition requires an extended period of due diligence, regulatory approvals, and post-acquisition integration. In 2009, after Fincantieri acquired Manitowoc Marine Group—which included Marinette Marine and smaller yards in Sturgeon Bay and Green Bay, Wisconsin—it took several years to achieve full operational integration (Kington, 2015).

Committee on Foreign Investment in the United States (CFIUS) review processes typically take 12–18 months, and conditional approvals limit operational integration (Roberts & Clark, 2025). Facility renovations, workforce training, and production process standardization require multiple years before meaningful productivity improvements materialize.

Technology transfer partnerships, particularly in the commercial sector, offer the most rapid implementation timeline but typically feature less dramatic ramp-ups of capital investment and output. A partnership dating to 2006 between General Dynamics-NASSCO's commercial division and DSEC, a division (at the time) of the South Korean marine engineering firm Daewoo Shipbuilding & Marine Engineering (DSME), began yielding productivity improvements within a year of inception (Heim & Tedesco, 2009, p. 9). Because such partnerships avoid ownership transition periods or CFIUS review (as well as ITAR compliance burdens if they exclusively focus on commercial vessels), they can commence relatively quickly.

Cost Implications

Direct acquisition of U.S. yards by foreign firms may require only limited upfront U.S. government expense, as capital investments typically come primarily from private foreign sources. However, foreign shipbuilders have sought U.S. government support for facility upgrades. Fincantieri paired its investment of approximately $500 million in Marinette Marine over the first 15 years of operation with $50 million each from the Navy and Wisconsin state government (Katz, 2023; Schumacher, 2022). Austal USA's recent expansion was supported by $288 million in private investment supplemented by $152 million in DPA Title III funding (Samora, 2024).

This public-private cost sharing reduces the fiscal burden on taxpayers compared to government-funded yard modernization. The South Korean government has pledged to invest as much as $150 billion in American shipbuilding as part of the bilateral "Make American Shipbuilding Great Again" (MASGA) initiative from 2025. The Trump administration's proposed maritime opportunity zones could spur additional private capital interest (J. Park et al., 2025; The White House, 2025a). However, indirect costs may arise from CFIUS reviews, compliance monitoring, and the suppression of alternative uses for waterfront real estate.

Non-equity partnerships like the NASSCO-DSEC collaboration impose lower upfront costs for both the government and private sector. NASSCO's partnership required limited facility investments (such as new paint-application systems) to achieve a reported 60% reduction in unspecified labor hour requirements for naval support vessels (White, 2011). The government's only unavoidable costs would be limited to regulatory oversight.

Bolstering U.S. Shipbuilding Capacity

The capacity creation potential of foreign acquisition and partnership is highly variable, dependent not only on the selected approach but also on implementation.

Direct acquisition has demonstrated the most dramatic instances of capacity creation in historical precedent. Fincantieri's acquisition of and investment in Marinette Marine transformed a small coastal vessel producer into a yard now capable of producing complex surface combatants including the Freedom-class LCS and Constellation-class frigates (Shelbourne, 2023). Austal USA rebuilt a similarly minor yard as a prime contractor for Independence-class LCS and eventually Virginia-class submarine modules—capabilities that would not have developed absent foreign investment and technology transfer (Austal USA, 2019).

The capacity creation mechanisms include physical infrastructure expansion (Fincantieri added approximately $500 million in facilities, including assembly buildings, ship lifts, and panel lines), workforce development (Marinette employment tripled), and technical capability enhancement through transfer of advanced manufacturing techniques such as modular construction and zone outfitting (Lockheed Martin, 2021). These improvements create enduring capacity that will benefit U.S. shipbuilding even if the foreign owner eventually sells the yard.

Commercial yard conversion through foreign investment could create substantial new naval capacity by revitalizing underutilized facilities. The U.S. has numerous marine sites with basic infrastructure (waterfront access, docks, cranes) that lack the specialized capabilities, infrastructure, and workforce needed for naval production. Foreign investment could provide capital and expertise to bridge this gap.

Technology transfer partnerships create capacity through knowledge diffusion rather than capital investment. NASSCO's reported 60% reduction in labor hours for naval support vessels represented real capacity creation, meaning the same workforce and facilities could now produce significantly more ships per year (White, 2011). However, this productivity-driven capacity expansion subsequently faltered in the face of reduced demand. Nonetheless, technology transfer partnerships create less capacity than direct acquisition because they lack the sustained capital commitment and operational control that enable transformative change.

Allied Industry Willingness and Ability

Allied shipbuilders have several reasons to favor investing in the United States. South Korean and Japanese shipbuilders must contend with labor shortages at home and competition with Chinese shipbuilders in the commercial market. Entering the United States would open a new market for these shipbuilders. Moreover, they may face pressure from their governments to cooperate as a way to promote country-to-country ties (Guevera, 2025; Oh & Cecire, 2025). The South Korean government in particular has emphasized its MASGA project as a way to manage and improve its broader relationship with the United States in light of tariff disputes and regional security challenges (Yu, 2025).

Given the rising U.S. shipbuilding budget, the large total addressable market of U.S. naval shipbuilding is also naturally attractive to foreign shipbuilders. The Trump administration's FY2027 budget request for shipbuilding was $65.8 billion, which is the largest shipbuilding budget proposal since 1962 when adjusted for inflation (Shelbourne & LaGrone, 2026).

The recent history of foreign shipbuilders purchasing or even partnering with U.S. yards is not entirely positive, and past foreign buyers have faced substantial economic challenges in the U.S. market. Entering the U.S. market is complicated for foreign shipyards because of compliance burdens and an uncertain demand signal from the U.S. government. Austal posted its worst financial results ever in 2003 because of lower than anticipated labor productivity and inexperience with security requirements (Austal Limited, 2003, p. 5). Even well-capitalized acquirers like Fincantieri have needed to invest approximately $500 million in capital improvements, process improvements, and workers in addition to the initial $120 million acquisition price of Marinette Marine in 2009 (Katz, 2023; Schumacher, 2022). ITAR restrictions impose significant additional costs by forcing international shipbuilders to establish entirely new, duplicative engineering teams.

Shipbuilders consistently emphasize that sustained capital expenditures require long-term, stable contracts.[1] Without assured demand, foreign firms are understandably reluctant to investment in a monopsonistic market with massive regulatory uncertainty, compliance costs, and workforce challenges.

Hanwha Ocean and HD Hyundai Heavy Industries, South Korea's two largest shipbuilders, have pursed divergent approaches to entering the U.S. market. Hanwha purchased Philly Shipyard, committing itself to making billions of dollars in capital investments and immense workforce development challenges (Ha, 2025). Hyundai has signed a more limited MOU with Huntington Ingalls Industries focused on joint research and development, MRO contracts, and bidding for next-generation logistics vessels (HII, 2025). Hanwha's strategy involves a higher degree of risk in the pursuit of a commanding portion in the U.S. Jones Act

[1] Interview with International Shipbuilder, June 23, 2025.

and defense markets, while Hyundai's is significantly more cautious, with more limited upside in the form of maintenance and module contracts.

U.S. Political and Regulatory Viability

Political support for foreign investment appears to be widespread. Former Navy Secretary Carlos Del Toro explicitly encouraged foreign investment in U.S. shipyards in 2024, praising Fincantieri and Austal for modernizing facilities while criticizing American shipbuilders for prioritizing stock buybacks over infrastructure investment (Lagrone, 2024). Navy Secretary John Phelan has stressed the need to imitate international best practices and praised Hanwha's investment in Philly Shipyard (Luckenbaugh, 2025). Local officials and congressional representatives in Alabama, Wisconsin, and Pennsylvania have consistently praised international shipbuilders for investing in their constituencies.

Not all actors in the United States offer full-throated support for bringing foreign shipbuilders into the U.S. industrial base. Some established U.S. naval shipbuilders view foreign investment as competitive threats that could win contracts that they want and undermine their business models. These concerns are not entirely unfounded. If foreign-owned yards capture a larger share of naval contracts due to superior efficiency or lower costs, purely domestic shipyards may face reduced order books and workforce instability.

Regulatory barriers present formidable obstacles even when political support exists. Executives from an international shipbuilder with a U.S. subsidiary have described how ITAR restrictions forced them to maintain parallel engineering teams across national lines, seek clearance to move employees between facilities, and obtain waivers for technical assistance—requirements that undermined the knowledge-sharing benefits that should flow from foreign ownership.[2]

The two approaches face different levels of regulation. Direct acquisition faces the strictest scrutiny, especially for naval yards. Investing in and converting a commercial yard to be capable of military shipbuilding (such as is the case in Hanwha Philly) may encounter somewhat less regulatory resistance if initially focused on auxiliary or support vessels. Warships are more complex assets with more sensitive technology on board, putting foreign ownership of yards under greater scrutiny. Commercial technology transfer partnerships without ownership changes face the fewest regulatory barriers; the NASSCO-Daewoo collaboration entirely avoided CFIUS review or extensive ITAR complications.

Summary

Fincantieri's transformation of Marinette Marine and Austal's development of Bender Repair demonstrate that allied investment and expertise can transform uneconomical or obsolete facilities, creating genuine capacity gains for U.S. shipbuilding (Lockheed Martin, 2021).

However, ITAR restrictions create challenges for foreign parent companies, from providing engineering support, sharing technical information, or moving personnel freely between their U.S. subsidiaries, and overseas facilities undermine the technology transfer benefits that represent the primary strategic rationale for accepting foreign ownership (Roberts & Clark, 2025). Foreign investment also cannot negate fundamental U.S. shipbuilding challenges such as design volatility, unnecessarily stringent naval requirements, single-source supply chains, and insufficient workforce development systems. Fincantieri's Constellation-class program experienced substantial delays and cost overruns after U.S. Navy modifications added

[2] Interview with U.S. Shipbuilder, April 1, 2025.

500 metric tons of additional outfitting to the European FREMM baseline, demonstrating that allied shipbuilders also struggle with dysfunctional procurement paradigms (Oakley, 2024).

For treaty allies such as Japan and South Korea, streamlined Technical Assistance Agreements (TAAs) and pre-approved personnel exchange programs could facilitate the knowledge flows that make foreign investment transformative rather than merely cosmetic. The success of the NASSCO-DSEC commercial collaboration illustrates the potential of tech-transfer partnerships unencumbered by the regulations and requirements associated with naval construction.

Foreign shipbuilders consistently emphasize that substantial investment in U.S. facilities requires confidence in long-term, stable demand.[3] The Navy's historically inconsistent procurement patterns—characterized by starts, stops, and dramatic shifts in ship class requirements—make investing in serving the U.S. naval market unattractive compared to commercial markets where contracts are more predictable. Multi-year block buys, realistic budgets, and credible production schedules that account for learning curves could make investing in the U.S. naval industrial base more appealing to allies.

Coproduction

A second approach for working with allies to increase ship production is through *coproduction.* This involves the United States working with an allied partner on shipbuilding, with production distributed across the nations, typically following one of two approaches. The first, traditional modular coproduction, involves foreign yards manufacturing self-contained, pre-outfitted ship modules to be transported to U.S. shipyards for final assembly and integration (Schank et al., 2016). The second, or "green hull," involves foreign shipyards constructing complete hull structures to be sent to U.S. yards for systems outfitting and completion.[4]

Proponents argue that modular coproduction could enable the United States to scale warship construction more rapidly by enabling parallel work across multiple facilities, resulting in time and cost savings (McDonald, 2025; UK Ministry of Defence, 2022; VU Marine, 2025). However, this pathway introduces learning curves in coordination and technical standards while facing political and regulatory barriers similar to the outright purchase of foreign ships (Papavizas, 2024). Advocates stress that green hull procurement is relatively straightforward to implement but acknowledge a learning curve regarding outfitting and highlight inefficiencies in producing ships this way.

Background

Allied nations have employed modular and green hull approaches for decades. These pathways are most often employed when: 1) the recipient nation requires new vessels urgently; 2) recipient nations lack the domestic capacity, expertise, or the financial resources to build complete ships themselves; and 3) the foreign partner provides design, engineering, and production expertise as well as hardware. In some commercial cases, modular production is used to outsource parts of the ship to countries with lower cost production.

Multiple examples exist of allied modular coproduction. In 2020, Germany selected Netherlands-based Damen Naval to design and manage the construction of its F126 Frigate program, aiming to harness Dutch expertise while splitting construction among at least three German shipyards—Peene-Werft, Kiel, and Blohm + Voss—with the goal of retaining 70% of the work in Germany. However, the program has suffered delays of at least three years and significant cost overruns (Hoffmann, 2025). German officials attributed the crisis to failures in

[3] Interview with International Shipbuilder, June 23, 2025.
[4] Interviews with International Shipbuilders, July 2025.

Damen's design and manufacturing software, while Dutch sources have argued that German micromanagement, including thousands of complex specifications and a monthlong approval process for minor subsystems, had itself compounded issues. As of April 2026, the role of lead contractor on the F126 project is in the process of transitioning to Rheinmetall, who acquired several major German shipbuilders in early 2026 (Pitel, 2026).

For the UK's Type 31 frigate program, Babcock has subcontracted hull block construction since 2023 for HMS Active (the class's second vessel) to Polska Grupa Zbrojeniowa. The blocks were manufactured in Poland and transported by barge to Scotland for integration. The deal was intended to control costs, circumvent British capacity bottlenecks, and prepare the Polish workforce for their own frigate program (Allison, 2023; Mustoe, 2023). However, outsourcing the modules has generated political controversy and resistance in Britain, including from organized labor (Allison, 2019).

There are also several examples of green hull cooperation. From 2007–2015, Navantia constructed the hulls for Australia's Canberra-class LHDs in Spain from the keel to the flight deck (approximately 85% of the vessel's structure), then transported them to Australia, where BAE Systems Australia installed the island superstructure, combat systems, and communications equipment (Navantia, 2018).

The coproduction of the Mistral-class LPDs split ship construction between three yards from 2003–2007. France's DCN (now Naval Group) handled the majority of systems and outfitting while Chantiers de l'Atlantique (another French yard) built the habitability modules and Stocznia Remontowa in Poland took on steelwork and hull construction to lower labor costs (Smallman et al., 2011).

The Polish Border Guard's current offshore patrol vessel program involves a green hull arrangement in which French shipbuilder Socarenam won the tender to build a 70-meter OPV for Poland, but hull construction was subcontracted to NavireTech in Poland. After a year of work, the partially equipped hull was launched in July 2022 and towed to Socarenam's shipyard in France, where final outfitting and equipment installation was completed. The program is expected to last from 2020–2030 (Peruzzi, 2023).

Pathway Evaluation

Ship Delivery Time

Previous efforts and interviews with shipbuilders suggest that the coproduction approach allows for eventual accelerated vessel delivery but involves a substantial upfront period to establish the necessary corporate, regulatory, and technical frameworks.

Once established, modular construction has the potential to accelerate individual ship construction. Commercial industry sources claim this approach can halve construction timelines via overlapping construction phases while boosting efficiency (VU Marine, 2025). Among other benefits, modular construction can occur in enclosed environments, avoiding weather-related delays that impact traditional shipyards. The United States' maritime industry is already familiar with modular construction, easing the implementation process (Ramponi, 2025). U.S. Virginia-class submarine producers have implemented aspects of modular shipbuilding, with General Dynamics Electric Boat and HII Newport News each building specialized elements and transporting supermodules to each other for a rotating final assembly location (Smallman et al., 2011, p. 50). The U.S. maritime industrial base is already making increased use of modular construction for surface ships, including Eastern Shipbuilding Group's work on DDG-51 Arleigh Burke modules for HII (Schuler, 2025).

These techniques offer the potential to reduce shipbuilding time requirements in the long run – especially if the U.S. Navy is willing to embrace new methods, including larger ship sections that offer greater access to workers. Implementing modular coproduction between U.S. and allied shipyards would require additional preparation before construction can commence. This would involve negotiation and signing government-to-government agreements establishing the legal framework as well as industry-to-industry agreements. The latter would need to cover the development of detailed interface specifications defining how modules from different yards will connect and the establishment of common quality assurance protocols.

Modular coproduction would also require upgrades to U.S. shipyard infrastructure to handle module receipt and assembly as well as the training of U.S. shipyard workers in integration techniques. These investments in U.S. infrastructure (discussed in detail below) will take time to plan and execute.

In contrast, the green hull approach's main stumbling blocks lie in the post-delivery installation of equipment and systems, which could require significant labor-intensive superstructure deconstruction and reconstruction.[5] Ships are not typically built "outside in," and either the correct cabling and interfaces for complicated systems have to be laid by the building yard – requiring significant information sharing efforts – or deconstruction and rework will be needed in the assembling yard (Birkler et al., 2015, p. 42). This approach runs counter to the best practice of advanced outfitting. While the green hull approach is likely simpler and quicker to get started, the long run benefits on per ship construction time may not be as high given rework requirements and process inefficiencies. Nevertheless, the outfitting of largely complete hulls is within the capabilities of the U.S. workforce and physical infrastructure, and the overall cooperation activity could begin relatively quickly.

Cost Implications

Making use of modular coproduction would require major infrastructure investments to upgrade U.S. shipyards to receive and assemble large modules from abroad, align standards with allied yards, cover transportation costs, and manage the inefficiencies of distributed production compared to one optimized facility.[6]

Although some U.S. shipyards have begun to adopt modern modular construction techniques, significant further investment would almost certainly be required.[7] This would dovetail with U.S. efforts to implement "nation as a shipyard" policies which leverage distributed modular construction across the United States and the recent U.S. Navy (2026) investment of $900 million (alongside $1.5 billion in private capital) to create the "Factory of the Future" in Cherokee, Alabama, to manufacture submarine components (O'Rourke, 2025, p. 24).

To make modular coproduction successful, U.S. shipyards would likely have to invest in facilities to accommodate module delivery and positioning. This would involve purchasing additional heavy-lift cranes, implementing precision measurement systems, and possibly revamping production flows.[8] Allied yards would need to ensure modules met the U.S. Navy's survivability and systems requirements. Extensive coordination would be needed to ensure modules fabricated abroad met American quality and safety standards (Cook et al., 2025).[9] Coordination costs would include government officials' time to align standards, as well as the pass-through costs of adjusting tooling and machining practices.

---

[5] Interview with International Shipbuilder, July 2025.
[6] Interview with International Shipbuilder, July 2025; interview with U.S. Shipbuilding Expert, May 23, 2025.
[7] Interview with U.S. Shipbuilding Expert, May 27, 2025.
[8] Interview with International Shipbuilder, July 2025; interview with U.S. Shipbuilding Expert, May 23, 2025.
[9] Interview with International Shipbuilder, July 2025; interview with U.S. Shipbuilding Expert, June 6, 2025.

Transporting large modules is expensive.[10] Construction across multiple locations involves one yard constructing and joining together several small blocks into superblocks, weatherproofing them for sea-borne transport, and then sending them to another yard for assembly. Typically barges or specialized heavy lift ships are needed to handle these superblocks, and the limited availability of such transportation can raise costs (Smallman et al., 2011, p. 24). A fully outfitted engine room module or accommodation block can weigh several hundred tons and require specialized heavy-lift vessels for transpacific transport (BreakBulk News, 2016). Modules must be protected from saltwater corrosion, mechanical stress, and shrinkage/expansion during transportation, adding cost and risk.[11]

Challenges to implemented allied modular coproduction include cross-shipyard communication across corporate lines, which is especially complicated by a lack of accountability when failures occur (Schank et al., 2011). Some of the difficulties coordinating between shipyards could be ameliorated if the ones overseas and in the United States were owned by the same overall entity.

The green hull approach might cost less to implement initially, but its relative inefficiency could lead to higher per-ship expense in the long run. Upfront capital and coordination costs would likely be lower because the approach keeps high-value, high-complexity, security-sensitive outfitting work in the United States. Transporting hulls is less sensitive than that of modules; Navantia's successful hull export from Spain to Australia demonstrates that a similar trans-Pacific arrangement should be possible.[12] However, building a ship in this manner is not a best practice. The significant reworking required to cut open parts of the ship (deck, piping, etc.) to emplace systems into the green hull would drive long-term cost increases, which must be balanced against lower up-front implementation costs.

Bolstering U.S. Shipbuilding Capacity

Modular coproduction's primary potential benefit to U.S. shipbuilding capacity would be through technology transfer and workforce skill development. By engaging with Korean and Japanese shipbuilders who employ advanced modular construction techniques, U.S. yards could acquire new methodologies that improve long-term productivity (Kuzminski & Schmiegel, 2025; Potter, 2024). While the United States has some experience in modular shipbuilding, there is still room to grow. Many sections of the U.S. shipbuilding industry already employ these methods, especially for nuclear submarines and some surface ships programs, such as the LPD-17 Flight I after Hurricane Katrina and increasingly Flight III Arleigh Burke destroyers (LaGrone & Shelbourne, 2025; O'Rourke, 2025, p. 25).

However, the modular coproduction pathway is not purely focused on expanding domestic production. Work performed in Korean and Japanese shipyards does not build U.S. capacity, except in the case where U.S. workers are brought over to learn skills working in foreign yards, such as in the Arctic Security Cutter program (Singsit, 2026). If U.S. yards change their workflow to only or primarily perform assembly of foreign-supplied blocks, U.S. yards could develop integration skills but likely would not markedly improve their overall ship construction capabilities.

The green hull approach may involve less domestic work than modularity construction, as many of the ships' blocks will have been joined in the foreign shipyard. Outfitting the hull structure after construction with high-end systems such as radar and weapons requires different skills than the advanced fabrication and pre-outfitting that support new ship construction.

[10] A 2011 RAND study of multi-yard cooperation efforts for warships noted transportation needs had implications on cost in all seven of their case studies (Smallman et al., 2011, p. 28).
[11] Interview with U.S. Shipbuilding Expert, May 23, 2025.
[12] Interview with International Shipbuilder, July 2025.

Modular integration demands precision engineering, complex project management, and sophisticated quality control—capabilities that could translate to improved overall shipbuilding prowess (VU Marine, 2025).

Allied Industry Willingness and Ability

Opinions differ among Japanese and Korean shipbuilders regarding the value and viability of green hull production for the U.S. market. Korean shipbuilders have voiced optimism that a green hull approach is more politically realistic than outright U.S. purchase of their vessels and could circumvent ITAR restrictions.[13] Korean builders have also noted that while the green hull approach would be less complex than modular cooperation, they have concerns about post-delivery systems integration and the need for intensive communication with U.S. prime contractors.[14] Japanese shipbuilders have raised concerns about the economics and quality control of the green hull approach, as well.[15]

Allied shipbuilders have expressed confidence that module production could decrease costs and timelines for U.S. ships, particularly if subcontractors are shared between the shipyards, creating commonality and efficiencies of scale from the start. For example, the Korean shipbuilder HD Hyundai Heavy Industries has signed an MOU with HII with one stated goal being to explore investing in distributed shipbuilding (HII, 2025).

Both Korean and Japanese yards have the capacity to carry out this form of cooperation. Modular production is central to how their yards function and is a core element of the cost-competitiveness that has driven their business success. Several South Korean shipbuilding firms, for example, already make extensive use of ship modules constructed abroad in yards in Vietnam and Singapore, amongst other locations, with final assembly in South Korea (Sung & Kim, 2025).

Similarly, constructing a green hull is not dissimilar from working to construct an entire ship – given differences in leaving access open to core systems that would need to be added later. Korean and Japanese yards have the facilities, workforce, and know-how to carry this out; however, it is worth noting that constructing only modules or green hulls could occupy their yards' productive capabilities that could otherwise possibly be deployed for potentially more profitable full-ship construction. This may reduce their willingness to cooperate – or drive-up expenses to compensate for their opportunity costs.

U.S. Political and Regulatory Viability

Both modular and green hull coproduction would likely encounter opposition from U.S. industry and labor, while contending with statutory restrictions that explicitly prohibit foreign construction of major vessel components.

Some U.S. shipbuilders, maritime unions, and aligned experts view modular or green hull coproduction as no different from buying complete foreign-built ships, which they have consistently opposed. This perspective is summarized in the words of one U.S. shipbuilding advocate who called coproduction with foreign shipyards a "disastrous mistake" that would involve "kicking American shipyard workers to the curb" (Paxton, 2024). Many political figures remain firmly behind maintaining and even strengthening protectionist restrictions (Tammy Baldwin, 2023).

Shipbuilders and their political backers oppose coproduction because it may threaten domestic shipyards' revenue while offering only limited opportunities for certain segments of the

[13] Interview with International Shipbuilder and International Defense Officials, June 30, 2025.
[14] Interview with International Shipbuilder, July 1, 2025.
[15] Interview with International Shipbuilder, July 1, 2025.

maritime industrial base. Final outfitting work, while valuable, is less remunerative than full ship construction, so if substantial portions of a vessel are fabricated in South Korea or Japan, U.S. yards lose out on the potential revenue and profit from labor-hours, change orders, and materials (Umbrex, n.d.). U.S. shipbuilders and their advocates also argue that coproduction might be a first step towards the United States purchasing foreign-built ships, which they oppose.

However, smaller U.S. yards and facilities that lack the capacity to bid on full-ship contracts might gain opportunities as module assembly or green-hull outfitting yards. (Naval Sea Systems Command, n.d.). Additionally, equipment and sustainment suppliers, particularly weapons and electronics system subcontractors, receive business regardless of where construction occurs and would benefit from increased volume (Titmuss, 2024).

Regulatory barriers compound political opposition. The Byrnes-Tollefson Amendment (10 U.S.C. § 8679) explicitly prohibits the construction of U.S. naval vessels, or "major component[s] of the hull or superstructure," in foreign shipyards. The statute includes a waiver provision, but annual defense appropriation acts, including the 2026 DOD Appropriations Act, consistently contain non-waivable provisions that prohibit the use of funds for constructing Navy ships at foreign shipyards (K. Park, 2025). The Coast Guard as a component of the Department of Homeland Security is not subject to this DOD-specific restriction on the use of the Shipbuilding and Conversion Funds in the Navy budget. This enabled the Trump administration to forge the icebreaker deal for the Arctic Security Cutter with Finnish and Canadian companies via presidential waiver.

Summary

Coproduction is a promising pathway for cooperation if policymakers wish to balance delivery speed with bolstering U.S. capacity. The Navy already encourages shipyards to distribute shipbuilding across the United States; subcontracting work for existing U.S. ship classes out to allied shipyards would accelerate the number of ships the United States can field across the next decade. It would necessitate expensive investments in U.S. shipyards to help receive and assemble these allied modules. But while some construction dollars would flow overseas, these investments could bolster U.S. shipbuilding capacity and competitiveness long term. The program could also serve as a way of accelerating U.S. adoption of the most advanced modular production methods.

Nonetheless, the fundamental challenge is that modular and green hull coproduction combine the disadvantages of both domestic and foreign production while capturing only some of the benefits. Both would be more expensive and time-consuming than building complete ships in Korea or Japan because of coordination challenges, transportation costs, and management overhead. Either approach would offer fewer work or profit opportunities for the domestic maritime industrial industry than purely U.S.-based construction. Both face much the same political and regulatory barriers as direct foreign ship purchases because U.S. industry and law regard foreign-built modules or green hulls as equivalent to foreign-built ships.

Green hull and modular coproduction approaches offer different upsides. The green hull approach is generally regarded as easier to implement by some international shipbuilders and some U.S. experts, though other international shipbuilders thought it may be difficult to implement and were concerned by quality control and interface compatibility.[16] It is possible that the assembly phase of this program could be carried out in less advanced shipyards in the

[16] Interview with International Shipbuilder July 1, 2025; interview with U.S. Shipbuilding Expert, May 23, 2025; interview with International Shipbuilder July 7, 2025.

domestic industrial base, enabling broader distribution of shipbuilding labor and dollars across the U.S. ecosystem – assuming that quality and technical difficulties can be overcome.

By contrast, modular coproduction offers greater potential for greater domestic participation and workforce skill training for the U.S. maritime industrial base. This pathway would be most viable under specific conditions, beginning with Congress not placing restrictions on shipbuilding funding in annual DOD appropriations acts. A program could be structured as a bridge with built-in transition mechanisms to secure political support. Initially, ships might involve 60–70% foreign module content, decreasing to 50% for mid-program ships, and ultimately achieving predominantly domestic construction. A viable program model might also include technology transfer and workforce development provisions, with allied shipbuilders required to train U.S. workers, share production methodologies, and support U.S. yards in developing modular production expertise.

Additionally, either method could be confined at first to specific ship classes. Support and auxiliary vessels may represent better candidates than surface warfare ships. They are similar to the ships that foreign shipyards have the greatest experience with and feature lower security risks, reduced complexity, and less political significance.

Purchase of Foreign Ships

Not all maritime nations have capacity to produce naval vessels within their country or to meet all their naval needs. Those nations often contract with shipyards in partner countries to procure ships and scale naval warfighting capability. Some countries may focus on building particular classes of ships while importing others. Earlier analysis suggests foreign acquisition typically occurs when domestic capacity is insufficient, unaffordable, or slow; to achieve cost-prohibitive or highly specialized capabilities; to balance efficiency and domestic participation with multi-yard contracts, post-delivery outfitting, or phased-in domestic participation; or to achieve technology transfer and workforce development (Price, 1984). In short, there are many reasons to import ships, and it is a common strategy to build a national navy using imported ships or mixing imported and domestically produced ships.

Although there are many examples of nations procuring ships from overseas, the U.S. Navy has limited recent experience with this approach and last procured foreign-built ships in the early 1980s from European shipyards (MARAD, 2024). These were fast sealift ships for Military Sealift Command rather than combatants.

To overcome challenges in the U.S. naval shipbuilding industrial base, numerous analysts have proposed procuring naval vessels built overseas – especially support ships (DiMascio, 2024; Montgomery, 2026; Wills, 2026). Former Congressional Research Service analyst Ronald O'Rourke (2026b) has argued that if the Navy needed to procure ships quickly, the only way to do this would be to buy foreign-made ships.

As with the other pathways, foreign ship purchase could take multiple forms, each with distinct tradeoffs between speed, cost, and fleet compatibility. The two potential approaches the United States could consider include 1) purchasing existing allied designs and 2) contracting with foreign yards to produce U.S. designs.[17]

[17] A third option for purchasing from U.S. allies would be joint development of new ship designs. While technically possible, this option is not considered here as the time, money, and intra-governmental effort required to create a new joint project rules this option out in the time scale and cost picture driving current cooperation needs. The need for two or more parties to agree when they have differing strategic needs, technical standards, political constituencies, and cost tolerances can cause roadblocks and inefficiencies. Recent U.S. efforts to adapt international designs

Pathway Evaluation

Ship Delivery Time

Allied shipyards could notionally produce and deliver ships relatively quickly provided there was clarity on the design and access to the necessary components in the supply chain.

South Korean and Japanese shipyards are accustomed to working quickly on shipbuilding projects, primarily commercial but also naval. Allied yards predominantly build commercial vessels, a deeply competitive market where success depends on delivering ships on time and at cost. Japanese and South Korean shipyards meet the challenge through attention to schedule and producibility, while investing in new production technologies (Jaquith, 2019). South Korea's acquisition practice is typically to fix requirements and limit change orders during the build process. Those lessons and best practices such as pre-outfitting of modules can apply to naval shipbuilding, even with different requirements. In practice, South Korea's adaptations of U.S. naval designs for their own fleet have consistently been delivered more quickly and consistently than their U.S. counterparts, though of course these ships are different – making specific estimates of time savings difficult (Bisht, 2021; Cha, 2024; Lee, 2021).

Procuring foreign-designed ships would be the fastest way to add ships to the U.S. naval fleet and would take advantage of existing production practices and supply chains. The shipbuilders would need to work with their U.S. Navy customer to understand any specific additional requirements. Further work would be necessary to incorporate U.S. naval systems on the vessels, and sensitive classified systems would likely need to be added in the United States.

Having foreign yards construct U.S. designs under license would require additional engagement to enable the allied shipbuilder to work with complex and often unique U.S. Navy requirements. Potential challenges include compliance, ensuring the correct processes are followed, and that security concerns are met. Foreign shipyards would have to follow the U.S. Navy's specific survivability specifications, systems integration protocols, and operational requirements.[18] At the basic production level, foreign yards typically use the metric system instead of the imperial system. Differences require either retooling or negotiation, although these delays are not insurmountable. Complex systems like the United States' AEGIS have been incorporated into allied ships by partner yards in the past (LaGrone, 2016).

Cost Implications

Japanese and South Korean naval ships are constructed in shipyards that build both commercial and naval vessels, where the push to be competitive in a commercial market increasingly dominated by Chinese ships forces builders to invest in increasing efficiencies and driving down costs. These investments, improvements, and management practices can be incorporated into naval production, where for example producibility is emphasized during the design phase.[19]

Procuring existing allied designs where the shipyard has already built ships of the specified design would be less expensive than purchasing a license-built U.S. design in terms of up-front purchase price. For example, comparing the cost of the FFG-62, DDG-51, and T-AOL

with the Constellation class have resulted in significant cost overruns (LaGrone & Shelbourne, 2024). Even within a single military, joint programs across services without identical and stable requirements can lead to cost growth, as the F-35 program showed (Lorell et al., 2013).

[18] Interview with U.S. Shipbuilder, October 23, 2025.

[19] This is not to imply that allied shipyards build commercial vessels and naval vessels interchangeably with the same workforces and facilities. Often, naval production occurs in a fenced-off section of a shipyard – and with a workforce comprised of the respective country's own nationals, which is not always the case in commercial construction. Nevertheless, some investment in commercial technology and infrastructure can benefit naval shipbuilding, not to mention the development of best practices in construction and management.

classes to the Korean Daegu, Sejong the Great, and Soyang-class vessels suggests that direct purchase of Korean designs could yield vessels 50–70% cheaper than their U.S. counterparts, even considering that the U.S. ships have different standards that also drive costs (Bisht, 2021; Cha, 2024; Lee, 2021). The Korean government notes that their Jeongjo the Great AEGIS-equipped destroyer was built in 38 months at a cost of around $1 billion, while the Arleigh Burke Flight-III has an approximately 4-year construction period and $1.4 billion dollar price.[20] However, life cycle logistics and MRO costs would be increased by the introduction of additional ship variants into the fleet, new parts, and different technical standards. Given that MRO typically accounts for a large proportion of vessel's lifetime cost, this could be a significant cost burden for this option in the long run (Maurer, 2024).

Producing U.S. designs rather than allied designs in allied yards would increase production costs for many reasons, and what follows is not an all-inclusive list. The U.S. Navy often evolves designs over the course of a build, especially for the first-in-class ship, which adds costs. South Korean yards, for example, typically require design maturity before they start laying the keel. These shipyards focus on "design for producibility," a term of art which means bringing ease-of-production considerations into the ship design process. As a result, their naval ships are often larger than comparable U.S. ships to enable easier worker access and build processes. Japanese ship construction practice also emphasizes minimal design changes during construction.[21]

U.S. Shipbuilding Capacity

Building ships abroad does not contribute to enhancing U.S. domestic shipbuilding capacity unless included as part of a hybrid strategy – discussed later – wherein foreign purchase was paired with a contractual requirement to invest in U.S. yards, as is commonly done in "offset" strategies.

If ship orders given to foreign yards were restricted to classes of ships that U.S. shipyards do not have the capacity or economic incentive to build, or if the Navy were to procure ships to the extent that U.S. shipyards' order books were overflowing, then sending orders abroad may not negatively impact the naval shipbuilding industrial base directly. However, this would still not contribute to the United States' capacity in the long run.

Allied Industry Willingness and Ability

Allied shipbuilders have indicated an interest in constructing ships for the U.S. Navy. In interviews, they have suggested that they could dedicate entire sections of their yards to U.S. naval construction if they received long term contracts.[22] Korean yards' success in building modified U.S. designs makes them confident they can deliver ships to the U.S. Navy at a substantially lower cost.[23] From their perspective, building their own designs would be easiest. Working with U.S. naval designs would be possible but challenging; along with different standards, foreign yards have highlighted the challenge of managing requirements creep given the allied yards' practice of constructing mature (i.e., fixed and unchanging) designs.

Moreover, for exporting nations, selling ships has the same benefits as other exports, including market expansion and support of the local economy, risk diversification, and economies of scale. The South Korean and Japanese governments are strong proponents of

[20] Slide deck provided by Ministry of National Defense, Republic of Korea.
[21] Interview with International Defense Officials, July 4, 2025.
[22] Interview with International Shipbuilder, September 24, 2025.
[23] Interview with International Shipbuilder, June 30, 2025; Interview with International Shipbuilder, July 2, 2025.

their builders selling to the United States as it would bolster their maritime industrial bases while simultaneously strengthening their defense partnerships with a crucial ally.

U.S. Political and Regulatory Viability

For nations that traditionally construct their own warships – including the United States – purchasing ships from overseas represents a change in strategy with significant barriers and internal resistance to overcome. Purchasing naval ships from partners and allies faces strong headwinds, with the most profound challenge based in U.S. law. The Byrnes-Tollefson amendment (10 U.S.C. § 8679) prohibits U.S. armed forces' vessels from being built in a foreign shipyard, though it can be waived by the president (Construction of Vessels in Foreign Shipyards: Prohibition, n.d.). Members of Congress have made consistent efforts to tighten these restrictions in annual NDAAs (Tammy Baldwin, 2023). While the 10 U.S.C. § 8679 restriction can be waived by the president, other restrictions in annual DOD appropriations acts on the use of Navy shipbuilding funds cannot be waived by the president acting alone. They would require the support of Congress (K. Park, 2025).

American shipbuilders have consistently opposed outsourcing naval vessel construction. They have claimed that they are far from full capacity and that current naval procurement barely supports their continued business health.[24] Unions such as the International Association of Machinists and Aerospace Workers have warned that that outsourcing would undermine national security (IAM Union, 2025).

These challenges, though difficult, are not insurmountable. The possibility of a presidential waiver in case of a national security emergency means that if there were no other alternative to growing the fleet, it would be possible to do so by buying foreign ships for the Navy – if Congress does not continue to place restrictions on naval funding.

Summary

Every pathway for cooperation presents distinct opportunities, and their feasibility depends on the willingness of U.S. leaders to overcome the challenges. The potential to leverage the direct procurement of foreign-built vessels from allied shipyards depends on strategic and political priorities. If U.S. policymakers' top priorities are speed, buying foreign ships is a strong option. However, if increasing domestic shipbuilding capacity is policymakers' first priority, procuring ships built in foreign yards will not yield tangible progress. Political and regulatory challenges make buying foreign ships difficult unless sufficient political will is mustered to change the current policy landscape.

Hybrid Approaches and Conclusions

This report does not seek to suggest a single optimal approach to address the shipbuilding challenge facing the United States, which include a series of interwoven, systemic issues in both government and industry that impede the industrial base's ability to build ships at the speed, scale, and cost the Navy needs. Each of the three aforementioned pathways—allied partnership with, investment in, or acquisition of U.S. shipyards; modular coproduction with foreign yards; and purchasing foreign ships—involves tradeoffs in terms of time, cost, and improvement of U.S. industrial capacity. These pathways also depend on both allied willingness and U.S. political and regulatory viability. Allied cooperation–regardless of the form it takes–will face challenges to implementation. Nonetheless, policymakers with distinct policy preferences can leverage allied capability to solve discreet problems, even if it does not immediately solve the entire problem set bedeviling U.S. naval shipbuilding.

[24] Interview with U.S. Shipbuilder, October 23, 2025.

If U.S. policymakers' focus were solely on rebuilding U.S. domestic capability, then encouraging allied investment in, partnership with, or acquisition of U.S. shipyards would be a key pathway – but it would likely not be a fast or cheap solution to deliver more ships into the U.S. fleet. It would also require some regulatory reform to enable easier communication and coordination between shipyards' foreign headquarters and their U.S. partners or subsidiaries.

If, instead, U.S. policymakers' top priority was delivering ships into the U.S. fleet as quickly as possible, then buying foreign designs off the shelf would present a more immediate option. But this will not bolster the capacity of the U.S. maritime industrial base in the long term. Although foreign ships may be cheaper than equivalent U.S. ships, they could end up costing more to sustain in the long run due to a lack of commonality. It would also require the expenditure of a large amount of political will or capital to change existing laws, at a minimum reforming the DOD appropriation acts' annual restrictions on shipbuilding funding to allow for foreign construction with a presidential waiver, as is the case with the Byrnes-Tollefson amendment (10 U.S.C. § 8679).

Modular coproduction could serve to both improve U.S. domestic capacity and increase ship deliveries to the fleet. It leverages allied capabilities to scale the production of U.S. ships and could reduce cost and ship construction time in the long run, although it would require a significant amount of startup time and capital. It could help bolster U.S. shipbuilding capacity via technology transfer and investments, but perhaps not to the same extent that focusing on allied partnership and investment might. And it would face strong political headwinds from U.S. domestic industry and would require political reforms potentially to the same extent as directly purchasing U.S. ships.

However, the different pathways to cooperation outlined in this report need not be implemented in isolation. In practice, cooperation will often entail combining two or more pathways to deliver ships more quickly, control costs, facilitate the buildup of domestic production capability, attain U.S. political support, or make cooperation more appealing to allied industry. Australia's planned procurement of the Mitsubishi Heavy Industries-designed Mogami frigate follows a hybrid approach. The lead ships will be built in Japan and subsequent ones in Australia (Australian Government Defense Ministers, 2025).

Another hybrid approach was recently employed by the United States with the Arctic Security Cutter (ASC) program, more commonly referred to as the Icebreaker Collaboration Effort (ICE) Pact.[25] Combining foreign construction, phased onshoring, and knowledge transfer via personnel exchange, the ASC program is a hybrid approach that uses foreign yards to build ships in the near term and to rebuild domestic shipbuilding capability via investment in and the sharing of expertise with the U.S. maritime industrial base. The ASC program involves both purchasing foreign ships and partnering with, investing in, and/or acquiring U.S. yards.[26] The United States was able to purchase foreign-built ships with a presidential waiver, without need

[25] In an example of cross-administration and bipartisan focus, the ICE Pact framework was established under the Biden administration in a November 2024 MOU, in which the United States, Canada and Finland agreed to collaborate on icebreaker construction. In April 2025, under the second Trump administration, the Coast Guard announced a Request for Information (RFI) for the Arctic Security Cutter, a program intended to supplement the delayed, over-budget polar security cutter program (O'Rourke, 2026a, p. 13).
[26] Contracts for up to 11 ASCs were divided between two contract teams. The first part of the contract went to Canadian shipbuilder Davie, which will construct two ASCs at its Helsinki Shipyard and subsequently three ASCs in Texas through its newly acquired U.S. subsidiary (formerly Gulf Copper; Naval News, 2026). In the second contract, the first two vessels will be built by Finnish shipbuilder Rauma in Finland using a design from Finland's Aker Arctic and Canada's Seaspan. Another four vessels will be constructed in the United States by Bollinger to the same design (U.S. Coast Guard News, 2025).

for congressional action, as the U.S. Coast Guard is part of DHS and thus subject to different procurement laws and regulations.[27]

The ASC program demonstrates that the United States has the capacity and will to engage in international shipbuilding cooperation, especially if done through a hybrid approach. Combining pathways in this way allows shipbuilding cooperation to perform well against many outcome metrics (such as building ships fast or building U.S. capacity) at once, while individual approaches require more difficult trade-offs to be made. Cooperation in hybrid models allows investment, partnership, coproduction, or even foreign purchase to all work hand-in-hand to make a deal both productive and politically viable.

That said, hybrid approaches do not offer a perfect solution. Hybrid approaches are more complex forms of cooperation, and basic defense industrial collaboration is already difficult to carry out. They also are a form of policy bundling, wherein every interest group may realize some objectives, but the solution may be subsequently less efficient (especially on cost) as a result. The choice between different hybrid approaches to collaboration, even beyond those in the ASC program, will depend on U.S. political tolerance for foreign ownership, the desired balance between speed and structural transformation, and foreign shipbuilders' capital access and risk tolerance.

Nevertheless, the ASC program demonstrates that when U.S. strategic needs are clear and pressing enough, it is possible to pursue innovative methods. Whether through an individual pathway or a combination of several, the United States should consider leveraging its allies and their shipbuilding industries to rebuild U.S. naval power.

References

Allison, G. (2019, January 4). *CSEU say government is 'slowly killing the shipbuilding industry.'* https://ukdefencejournal.org.uk/cseu-say-government-is-slowly-killing-the-shipbuilding-industry/

Allison, G. (2023, April 9). *Poland building part of a British frigate – why?* https://ukdefencejournal.org.uk/poland-building-part-of-a-british-frigate-why/

Austal Limited. (2003). *Concise report 2003*. https://investor.austal.com/static-files/402bad07-95cf-48ce-993e-89e4a8d2d38e

Austal USA. (2019, December 13). *Austal USA celebrates 20 years.* https://www.austalusa.com

Australian Government Defense Ministers. (2025, August 5). *Mogami-class frigate selected for the Navy's new general purpose frigates*. https://www.minister.defence.gov.au/media-releases/2025-08-05/mogami-class-frigate-selected-navys-new-general-purpose-frigates

Birkler, J., Schank, J. F., Arena, M. V., Keating, E. G., Predd, J. B., Black, J., Danescu, I. E., Jenkins, D., Kallimani, J. G., Lee, G. T., Lough, R., Murphy, R., Nicholls, D., Persi Paoli, G., Peetz, D., Perkinson, B., Sollinger, J. M., Tierney, S., & Younossi, O. (2015). *Australia's naval shipbuilding enterprise: Preparing for the 21st century*. https://www.rand.org/pubs/research_reports/RR1093.html

Bisht, I. S. (2021, August 26). South Korea approves $3.07 billion frigate program. *The Defense Post*. https://thedefensepost.com/2021/08/26/south-korea-frigate-program/

[27] President Trump issued a presidential memorandum authorizing construction of up to four Arctic Security Cutters abroad under 14 U.S.C. 1151(b), the Coast Guard version of 10 U.S.C. § 8679, which prohibits foreign construction unless waived by the president (The White House, 2025b). While Navy shipbuilding funds are restricted from use for overseas procurement via DOD authorization act provisions, the U.S. Coast Guard under DHS does not face these unwaivable congressional restrictions and therefore may acquire foreign-built vessels with a presidential waiver (Construction of Vessels in Foreign Shipyards: Prohibition, n.d.).

BreakBulk News. (2016, December 9). *Transporting modules a much bigger task than just crossing Pacific*. https://breakbulk.news/transporting-modules-much-bigger-task-just-crossing-pacific/

Carroll, H. H., & Cook, C. R. (2025). *Identifying pathways for U.S. shipbuilding cooperation with Northeast Asian allies*. Center for Strategic and International Studies. https://www.csis.org/analysis/identifying-pathways-us-shipbuilding-cooperation-northeast-asian-allies

Cha, E. (2024, August 7). *Hanwha Ocean to build second Soyang-class fast combat support ship (AOE-II)*. Naval News. https://www.navalnews.com/naval-news/2024/08/hanwha-ocean-to-build-second-soyang-class-fast-combat-support-ship-aoe-ii/

Construction of Vessels in Foreign Shipyards: Prohibition, 10 U.S.C. § 8679 (n.d.). Retrieved October 31, 2025, from https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title10-section8679&num=0&edition=prelim

Cook, C. R., Edel, C., Paik, K., Sanders, G., Carroll, H. H., Buda, K., & Augé, J. (2025). *Enhancing defense industrial cooperation between Australia and the United States*. https://www.csis.org/analysis/enhancing-defense-industrial-cooperation-between-australia-and-united-states

Daniels, S. P., Carroll, H. H., Cook, C. R., Buntin, O., & O'Rourke, S. (2025). *Outlining the challenges to U.S. naval shipbuilding*. https://www.csis.org/analysis/outlining-challenges-us-naval-shipbuilding

DiMascio, B. T. (2024, October 22). *Foreign shipyards can help the U.S. Navy build its fleet*. U.S. Naval Institute. https://www.usni.org/magazines/proceedings/2024/october/foreign-shipyards-can-help-us-navy-build-its-fleet

Guevera, J. (2025, April 17). *Shipbuilding and the future of U.S. sea power: The role of South Korea and Japan in the rejuvenation of the American fleet*. Center for Maritime Strategy. https://centerformaritimestrategy.org/publications/shipbuilding-and-the-future-of-u-s-sea-power-the-role-of-south-korea-and-japan-in-the-rejuvenation-of-the-american-fleet/

Ha, S.-Y. (2025, July 7). Trump's efforts to revive U.S. shipbuilding slowed by inconsistent policies, leadership gaps, funding shortfalls. *Korea JoonAng Daily*. https://koreajoongangdaily.joins.com/news/2025-07-07/business/industry/Trumps-efforts-to-revive-US-shipbuilding-slowed-by-inconsistent-policies-leadership-gaps-funding-shortfalls/2345718

Heim, A., & Tedesco, M. (2009, July 30). *A shipbuilder's assessment of America's marine highways*. General Dynamics NASSCO. https://www.nassco.com/pdfs/Shipbuilder-Assessment-American-Marine-Highway-NASSCO.pdf

HII. (2025, October 26). *HD Hyundai Heavy Industries and HII execute memorandum of agreement to collaborate on distributed shipbuilding and pursue teaming on auxiliary and commercial vessels*. https://hii.com/news/hd-hyundai-heavy-industries-and-hii-execute-memorandum-of-agreement-to-collaborate-on-distributed-shipbuilding-and-pursue-teaming-on-auxiliary-and-commercial-vessels/

Hoffmann, L. (2025, July 16). *First calls from German MP for cancellation of F126 frigate programme*. Naval News. https://www.navalnews.com/naval-news/2025/07/first-calls-from-german-mp-for-cancellation-of-f126-frigate-programme/

IAM Union. (2025, January 7). *IAM Union continues to champion expansion of domestic U.S. Navy shipbuilding*. https://www.goiam.org/news/imail/iam-union-continues-to-champion-expansion-of-domestic-u-s-navy-shipbuilding/

Jaquith, P. E. (2019). Asian vs. U.S. warship design, production engineering, and construction practice. *Naval Engineers Journal*, *131*(4), 55–58.

Katz, J. (2023, April 19). *Fincantieri finishing $300M shipyard renovations, a big bet on the US Navy's frigate plans*. Breaking Defense. https://breakingdefense.com/2023/04/fincantieri-finishing-300m-shipyard-renovations-a-big-bet-on-the-us-navys-frigate-plans/

Kington, T. (2015, May 16). Naval work booming at Fincantieri. *Defense News*. https://www.defensenews.com/naval/2015/05/16/naval-work-booming-at-fincantieri/
Kuzminski, K., & Schmiegel, L. (2025, July 3). *A workforce strategy for America's shipbuilding future*. War on the Rocks. https://warontherocks.com/2025/07/a-workforce-strategy-for-americas-shipbuilding-future/
LaGrone, S. (2016, September 6). *New South Korean destroyers to have BMD capability*. USNI News. https://news.usni.org/2016/09/06/new-south-korean-destroyers-ballistic-missile-defense-capability
Lagrone, S. (2024, March 7). *SECNAV Del Toro tells U.S. shipyards 'invest more', encourages foreign investment*. USNI News. https://news.usni.org/2024/03/07/secnav-del-toro-tells-u-s-shipyards-invest-more-encourages-foreign-investment
LaGrone, S., & Shelbourne, M. (2024). *Constellation frigate delivery delayed 3 years, says Navy*. USNI News. https://news.usni.org/2024/04/02/constellation-frigate-delivery-delayed-3-years-says-navy
LaGrone, S., & Shelbourne, M. (2025, September 12). *U.S. naval shipyards accelerating outsourcing for new construction programs*. USNI News. https://news.usni.org/2025/09/12/u-s-naval-shipyards-accelerating-outsourcing-for-new-construction-programs
Lee, D. (2021, November 12). *South Korea's HHI wins new KDX III Batch II destroyer contract*. Naval News. https://www.navalnews.com/naval-news/2021/11/south-koreas-hhi-wins-new-kdx-iii-batch-ii-destroyer-contract/
Lockheed Martin. (2021). *Shipyard transformation through investment* [Press release]. https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/lcs/Shipyard_Brochure.pdf
Lorell, M. A., Kennedy, M., Leonard, R. S., Munson, K., Abramzon, S., An, D. L., & Guffey, R. A. (2013). *Do joint fighter programs save money?* https://www.rand.org/pubs/monographs/MG1225.html
Luckenbaugh, J. (2025, June 2). Navy taking shipbuilding lessons from Indo-Pacific allies. *National Defense Magazine*. https://www.nationaldefensemagazine.org/articles/2025/6/2/just-in-navy-taking-shipbuilding-lessons-from-indopacific-allies
MARAD. (2024, September 27). *National register eligibility assessment vessel: Antares (AKR-294)*. https://vesselhistory.marad.dot.gov/documents/68f6e020-0b5b-48a1-aff5-8e4f2ebdc45d.pdf
Maurer, D. (2024). *Weapon system sustainment: DOD identified operating and support cost growth but needs to improve the consistency and completeness of information to Congress* (GAO-24-107378). U.S. GAO. https://www.gao.gov/products/gao-24-107378
McDonald, J. (2025, March 19). *Modular construction is key for Navy's shipbuilding goals*. GovCIO Media & Research. https://govciomedia.com/modular-construction-is-key-for-navys-shipbuilding-goals/
Montgomery, M. (2026, February 15). Trump could be missing the opportunity to rebuild the Navy efficiently and quickly. *New York Post.* https://nypost.com/2026/02/15/opinion/trump-could-be-missing-the-opportunity-to-rebuild-the-navy-efficiently-and-quickly/
Mustoe, H. (2023, April 9). Britain's new warships to be partly built in Poland. *The Telegraph*. https://www.telegraph.co.uk/business/2023/04/09/britains-new-frigates-built-in-poland/
Naval News. (2026, February 11). *Davie Defense awarded USCG contract to build five Arctic security cutters*. https://www.navalnews.com/naval-news/2026/02/davie-defense-awarded-uscg-contract-to-build-five-arctic-security-cutters/
Naval Sea Systems Command. (n.d.). *Navy partners with private industry to grow submarine industrial base capacity*. Retrieved November 10, 2025, from

https://www.navsea.navy.mil/Media/News/Article-View/Article/3913023/navy-partners-with-private-industry-to-grow-submarine-industrial-base-capacity/

Navantia. (2018, June 26). *HMAS Canberra handed over to the Royal Australian Navy*. https://www.navantia.es/en/news/press-releases/hmas-canberra-handed-over-to-the-royal-australian-navy/

Oakley, S. S. (2024, May 29). *Navy frigate: Unstable design has stalled construction and compromised delivery schedules*. U.S. GAO. https://www.gao.gov/products/gao-24-106546

Oh, M., & Cecire, M. (2025). *Why the United States, South Korea and Japan must cooperate on shipbuilding*. RAND. https://www.rand.org/pubs/commentary/2025/05/why-the-united-states-south-korea-and-japan-must-cooperate.html

O'Rourke, R. (2025). *Navy force structure and shipbuilding plans: Background and issues for Congress*. Congressional Research Service. https://sgp.fas.org/crs/weapons/RL32665.pdf

O'Rourke, R. (2026a). *Coast Guard polar security cutter (PSC) and Arctic security cutter (ASC) icebreaker programs: Background and issues for Congress* (CRS Report No. RL34391). Congressional Research Service. https://www.congress.gov/crs-product/RL34391

O'Rourke, R. (2026b, February 11). *Building the fleet to meet the mission* [Public comments on panel]. WEST 2026. https://www.youtube.com/watch?v=2wkCBwb5lXI

Papavizas, C. (2024, May 7). *When is a vessel built in America "U.S. built."* MaritimeFedWatch, Winston & Strawn. https://www.winston.com/en/blogs-and-podcasts/maritime-fedwatch/when-is-a-vessel-built-in-america-us-built

Park, J., Singh, K., & Singh, K. (2025, November 14). *US, South Korea unveil details on shipbuilding investment*. Reuters. https://www.reuters.com/world/asia-pacific/us-south-korea-release-details-deal-including-korean-investment-shipbuilding-2025-11-14/

Park, K. (2025, September 30). *U.S. defense bill upholds overseas shipbuilding ban, challenging MASGA*. The Chosun Daily. https://www.chosun.com/english/world-en/2025/10/01/GNSOKZHDOFGA7HA6OZIBGMMACI/

Paxton, M. (2024, August 5). *Outsourcing the US shipyard industrial base will outsource American sovereignty*. National Defense Transportation Association. https://www.ndtahq.com/outsourcing-the-us-shipyard-industrial-base-will-outsource-american-sovereignty/

Peruzzi, L. (2023, May 31). *Offshore patrol vessels and shipbuilders in Europe*. ESD. https://euro-sd.com/2023/05/articles/31365/offshore-patrol-vessels-and-shipbuilders-in-europe/

Pitel, L. (2026, April 7). Europe's rearmament meets reality: The story of a failed frigate project. *Financial Times*. https://www.ft.com/content/124c9dfc-18da-49fa-aab5-6389dce833ae?syn-25a6b1a6=1

Potter, B. (2024, May 3). *Lessons from shipbuilding productivity—Part I*. https://www.construction-physics.com/p/lessons-from-shipbuilding-productivity

Price, R. A. (1984). *Manufacturing technology for shipbuilding* (NSRP #0140; Shipbuilding Technology Transfer). U.S. Department of Transportation Maritime Administration. https://doi.org/10.21236/ADA443660

Ramponi, D. (2025, September 22). *Breaking ships into blocks: How modular construction is changing shipbuilding*. ShippingKnowHow. https://www.shippingknowhow.net/post/breaking-ships-into-blocks-how-modular-construction-is-changing-shipbuilding

Roberts, M., & Clark, B. (2025). *Shoring up the foundation: Affordable approaches to improve US and allied shipbuilding and ship repair*. Hudson Institute. https://www.hudson.org/supply-chains/shoring-foundation-affordable-approaches-improve-us-allied-shipbuilding-ship-repair-bryan-clark-michael-roberts

Samora, S. (2024, July 21). *Austal USA spending $288M to expand Alabama ship plant*. Manufacturing Dive. https://www.manufacturingdive.com/news/austal-usa-288-million-ship-plant-expansion-mobile-alabama-us-navy/722779/

Schank, J. F., Arena, M. V., Ip, C., Kamarck, K. N., Murphy, R. E., Lacroix, F. W., & Lee, G. T. (2011). *Lessons from Australia's Collins submarine program* (Vol. IV; Learning from Experience). RAND Corporation.

Schank, J. F., Savitz, S., Munson, K., Perkinson, B., McGee, J., & Sollinger, J. M. (2016). *Designing adaptable ships: Modularity and flexibility in future ship designs*. https://www.rand.org/pubs/research_reports/RR696.html

Schuler, M. (2025, September 16). *Eastern Shipbuilding partners with HII to boost U.S. Navy destroyer production*. gCaptain. https://gcaptain.com/eastern-shipbuilding-partners-with-hii-to-boost-u-s-navy-destroyer-production/

Schumacher, K. (2022, May 12). *Fincantieri Marinette Marine cuts the ribbon on a new building for new Navy ships*. WBAY. https://www.wbay.com/2022/05/12/fincantieri-marinette-marine-cuts-ribbon-new-building-new-navy-ships/

Shelbourne, M. (2023, April 25). *Major shift comes to Fincantieri Marinette Marine as Freedom LCS line ends*. USNI News. https://news.usni.org/2023/04/25/major-shift-comes-to-fincantieri-marinette-marine-as-freedom-lcs-line-ends

Shelbourne, M., & LaGrone, S. (2026, April 3). *Pentagon's new $65.8B shipbuilding ask is highest since 1962*. USNI News. https://news.usni.org/2026/04/03/pentagons-new-65-8b-shipbuilding-request-is-highest-since-1962

Silva, F. (2024). *Governmental influence on shipbuilding* [Master's thesis, Naval Postgraduate School]. https://dair.nps.edu/handle/123456789/5280

Singsit, J. (2026, February 12). *US Coast Guard selects Davie Defense for five Arctic security cutters*. Naval Technology. https://www.naval-technology.com/news/uscg-asc-davie-defense/

Smallman, L., Tang, H., Schank, J. F., & Pezard, S. (2011). *Shared modular build of warships: How a shared build can support future shipbuilding*. RAND. https://www.rand.org/pubs/technical_reports/TR852.html

Sung, S.-H., & Kim, J.-W. (2025, April 14). *With domestic dockyards fully booked, Korean shipbuilders turn overseas*. KED Global. https://www.kedglobal.com/shipping-shipbuilding/newsView/ked202504140008

Tammy Baldwin. (2023, July 21). *Baldwin strengthens Buy America requirements for Navy ships in must-pass Senate defense bill*. https://www.baldwin.senate.gov/news/press-releases/baldwin-strengthens-buy-america-requirements-for-navy-ships-in-must-pass-senate-defense-bill

Titmuss, P. (2024, April 5). *7 Naval contractors serving the U.S. government*. ExecutiveGov. https://www.executivegov.com/articles/naval-contractors-serving-us-government

UK Ministry of Defence. (2022). *Maritime modularity concept*. https://assets.publishing.service.gov.uk/media/63dcece38fa8f57fc8061089/Maritime_Modularity_Concept.pdf

Umbrex. (n.d.). *How the defense contractor industry works*. Retrieved November 10, 2025, from https://umbrex.com/resources/how-industries-work/aerospace-defense/how-the-defense-contractors-manufacturing-industry-works/

U.S. Coast Guard News. (2025, December 29). *Coast Guard awards contracts to build Arctic security cutters*. https://www.news.uscg.mil/Press-Releases/Article/4368196/coast-guard-awards-contracts-to-build-arctic-security-cutters/

U.S. Navy. (2026, March 20). *Advanced shipbuilding "factory of the future" opens in Alabama*. https://www.navy.mil/Press-Office/Press-Releases/display-pressreleases/Article/4439992/advanced-shipbuilding-factory-of-the-future-opens-in-alabama/

VU Marine. (2025, June 21). *Modular shipbuilding: A game changer in modern marine manufacturing*. https://vumarine.com/modular-shipbuilding-a-game-changer-in-modern-marine-manufacturing/

White, R. D. (2011, July 3). Full steam ahead for Nassco shipyard in San Diego. *Los Angeles Times*. https://www.latimes.com/business/la-xpm-2011-jul-03-la-fi-made-in-california-shipyard-20110703-story.html

The White House. (2025a, April 9). *Restoring America's maritime dominance*. https://www.whitehouse.gov/presidential-actions/2025/04/restoring-americas-maritime-dominance/

The White House. (2025b, October 8). *Memorandum on construction of Arctic security cutters*. https://www.presidency.ucsb.edu/documents/memorandum-construction-arctic-security-cutters

Wills, S. (2026, March 11). 'Let foreign yards build U.S. Navy auxiliary and service ships now.' *Seapower*. https://seapowermagazine.org/let-foreign-yards-build-u-s-navy-auxiliary-and-service-ships-now/

Yu, J. (2025, August 14). *How South Korea's 'MASGA' proposal could reshape US shipbuilding*. The Diplomat. https://thediplomat.com/2025/08/how-south-koreas-masga-proposal-could-reshape-us-shipbuilding/

Shipbuilding Procurement: An International Analysis of Source Selection Processes

Dr. Rene G. Rendon—is a professor of acquisition at the Naval Postgraduate School (NPS). A retired Air Force contracting officer, he served as a warranted contracting officer for major weapons system programs such as the Peacekeeper ICBM, F-22 Advanced Tactical Fighter, Space-Based Infrared Satellite program, and the Evolved Expendable Launch Vehicle program. He also served as a contracting squadron commander for an Air Force pilot training base. He was presented with the Air Force Outstanding Officer in Contracting Award, and he has received the NPS Hamming Teaching Excellence Award and the Graduate School of Business and Public Policy Research Excellence Award. [rgrendon@nps.edu]

Abstract

Naval shipbuilding procurement accounts for a significant portion of discretionary budgets worldwide. With high initial costs and long-term sustainment expenses, an effective source selection strategy is crucial in the contracting process for these critical weapon systems. This research examines the shipbuilding procurement source selection strategies of the United States, Egypt, Japan, and Taiwan, comparing these countries' source selection basis for award, source selection team composition, and proposal evaluation criteria. Using government acquisition regulations, laws, and publicly available solicitation/tender data, a comparative analysis is conducted to identify opportunities for improvement in the contract source selection process.

Introduction

Naval warships and support vessels are indispensable to the United States' ability to project power, maintain freedom of navigation, deter adversaries, and respond to crises worldwide. The U.S. Navy's maritime mission depends on a capable, ready, and technologically advanced fleet. However, despite significant and sustained investment in shipbuilding, the Navy has experienced persistent acquisition challenges that have limited fleet growth and readiness outcomes. Congressional testimony has highlighted that while the Navy's shipbuilding budget has increased substantially over the past two decades, systemic acquisition challenges continue to result in cost growth, schedule delays, and performance shortfalls (GAO, 2025a). These challenges underscore the strategic importance of improving shipbuilding procurement and contract management, specifically source selection processes, to ensure that delivered ships meet operational requirements.

The U.S. shipbuilding industrial base plays a central role in achieving cost, schedule and performance objectives. Private shipyard contractors and their subcontractors are responsible for constructing highly complex naval platforms within demanding cost and schedule constraints. Yet the industrial base faces structural limitations, including workforce shortages, aging infrastructure, and uneven workload distribution. The GAO has reported that the U.S. Navy lacks a fully integrated, strategic approach to private-sector shipyard investment, which introduces risk to planned fleet expansion goals (GAO, 2025b). Without a stable and capable industrial base, even well-designed acquisition strategies may fail to deliver ships on budget, on time, and in compliance with performance specifications.

The Navy shipbuilding contracting process serves as the formal mechanism linking Navy requirements with industry capability. Governed by the Federal Acquisition Regulation (FAR) and Department of Defense (DoD) FAR Supplement, this process is designed to promote competition, ensure fairness, and achieve best value for the government (FAR, 2026). Nevertheless, oversight findings indicate recurring weaknesses in the contracting life cycle pre-award, award, and post award phases. For example, the GAO has found that shipbuilding programs often begin construction with incomplete or unstable designs, contributing significantly to cost escalation and delayed delivery (GAO, 2024). Such systemic acquisition practices elevate the importance of rigorous source selection processes during the contracting life cycle award phase.

At the core of DoD contract management is the source selection process, which includes evaluating proposals, conducting cost or price analysis, negotiating contract terms, and selecting the contractor for award (NCMA, 2022). Decisions made in planning for source selection, as well as during source selection, have enduring implications for contract performance and ship delivery outcomes. Weaknesses in proposal evaluation or insufficient negotiation rigor can propagate risk into production, sustainment, and life cycle cost management. Auditability theory establishes the importance of capable processes, as well as competent people and effective internal controls for organizational success (Rendon & Rendon, 2015). Additionally, oversight bodies have emphasized that disciplined contracting practices are essential to achieving fair and reasonable pricing and effective contract performance (DoD OIG, 2025). Accordingly, understanding the Navy's ship contract source selection process is critical to evaluating broader acquisition success and fleet readiness outcomes.

This research examines the shipbuilding procurement source selection approaches of the United States, Egypt, Japan, and Taiwan, comparing these countries' source selection basis for award, source selection team composition, and proposal evaluation criteria. Using government acquisition regulations, laws, and publicly available solicitation/tender data, a comparative analysis is conducted to identify opportunities for improvement in the contract source selection process. While all four countries share similarities in their source selection process, their source selection basis for award, team composition, and proposal criteria differ in certain respects. Recommendations based on our findings are provided to improve efficiency and transparency in contract source selections, and long-term sustainability across multiple nations.

Background

The U.S. Office of the Secretary of the Navy released the FY2025 30-year shipbuilding plan, which calls for 381 crewed ships and 134 large, unmanned surface and underwater vessels as part of the optimal mix of ships in the U.S. arsenal (O'Rourke, 2024). The U.S. Navy maintains 296 battle force ships in its inventory, showing a significant disparity between the current and desired end state.

Egypt's tension in the Mediterranean Sea, geopolitical events, disputes over maritime borders, territorial waters, and gas reserves need a powerful Egyptian naval presence. Moreover, the Egyptian location at the intersection of the Red Sea, the Mediterranean Sea, and the Suez Canal is regarded as one of the most critical maritime corridors in the world. For these reasons, the Egyptian Ministry of Defense (EMOD) provides strategy that supports Navy forces to secure the country's interests and national security (EMOD, 2010).

In Japan, according to the National Defense Strategy, the Defense Buildup Program was announced in December 2022, which set out the level of defense capabilities that Japan should possess, the total expected cost of achieving this, and the quantity of major equipment to procure (Japanese Ministry of Defense [JMOD], 2024). According to this Defense Buildup Program, the plan is to procure two Aegis System-equipped vessels, 12 destroyers, five submarines, and 10 patrol vessels (PV) by March 2028 (JMOD, 2024).

In 2016, the Taiwan Ministry of National Defense (MND) announced 12 shipbuilding programs with a timeline from 2017 to 2040 at roughly $14.7 billion to enhance military power and support the self-reliant military establishment (Minnick, 2016). According to the National Defense Report from the Republic of China (ROC) in 2021 (Ministry of National Defense [MND], Taiwan, 2021), Taiwan's Navy had started shipbuilding construction for the Indigenous Defensive Submarine (IDS) program, High-Speed Minelayers, a new Landing Platform Dock (LPD), High-Performance Vessel (HPV) follow-up ships, as well as a new rescue and salvage ship. These plans demonstrated that Taiwan's shipbuilding procurement is committed to indigenous production instead of foreign ship purchases (Global Taiwan Institute, 2022). Thus, with increasing demands on indigenous shipbuilding in Taiwan, the procurement strategy for shipbuilding would be critical to select the most advantageous tenders and maximize the overall benefit for Taiwan's Navy.

Each country, driven by its unique motivations, national doctrines, and geopolitical climates, recognizes the strategic importance of naval procurement to secure national interests. The United States, Egypt, Japan, and Taiwan each demonstrate distinct objectives but share a common emphasis on bolstering naval forces to ensure maritime security. As allies and strategic partners, these nations might benefit from standardized or aligned procurement practices that potentially create efficiencies and identify areas for improvement. The purpose of this research is to analyze the source selection approaches of each country to assess whether adopting common source selection policies and practices could enhance their procurement efforts.

Problem Statement

The problem that motivated this research is that each country's source selection approaches related to shipbuilding are different in policy and practice. These disparities complicate the general understanding and potential international standardization in this field. By examining the source selection approaches used by other countries, especially in areas such as shipbuilding, nations can gain insights that may enhance their approaches to source selection and bridge a gap in cross-border cooperation.

Purpose Statement

In the global defense industry, procurement strategies play a central role in determining the overall success of maritime projects. The United States, Egypt, Japan, and Taiwan each employ distinct approaches to source selection in shipbuilding procurement, reflecting differences in priorities and regulatory environments. Despite this industry's importance, limited comparative research analyzes and compares source selection approaches among these nations. This apparent lack of analysis creates a significant gap in understanding how these countries address common challenges in contract management processes.

This research aims to clarify the differences in source selection approaches based on comparative analysis of the source selection evaluation processes, source selection team composition, and proposal evaluation criteria of the U.S. Navy, the Egyptian Navy, the Japan Maritime Self-Defense Force (JMSDF), and the Taiwan Navy. By drawing a meaningful connection between the different source selection approaches, this study aims to identify potential implications and areas for improvement and to contribute to the academic literature promoting standardization in international policy that could streamline further collaborative or individual efforts.

Research Questions

This research addresses the following questions:

1. How does the source selection approach (basis for award, source selection team composition, and proposal evaluation criteria) differ among the U.S. Navy, Egyptian Navy, JMSDF, and Taiwan Navy?
2. Based on the comparison and analysis, what implications for process improvement could be presented to the U.S. Navy, Egyptian Navy, JMSDF, and Taiwanese Navy?

Methodology

The methodology used in this research included a comprehensive review of the laws, regulations, and guidelines related to each country's contracting process, compared to the Contract Management Standard contract life cycle phases and related activities. Next, recent, real-world shipbuilding procurement solicitations, consisting of Request for Proposals (RFPs) and Request for Tenders (RFTs) from the United States, Egypt, Japan, and Taiwan, were analyzed to identify the source selection basis for award, source selection team composition, and proposal evaluation factors used in those shipbuilding procurements. Finally, the implications of the findings are discussed in terms of innovation, and flexibility in shipbuilding procurement. This paper is based on the graduate theses by Si-yun Yang (Yang, 2023) and Shinya Matsuda, Andrew C. Peters, and Amr A. Aboutaleb (Matsuda et al., 2025).

Findings

The research findings are discussed in terms of source selection basis for award approach, source selection team composition, and proposal evaluation factors for each country as reflected below.

Source Selection Basis for Award

Of the 30 United States shipbuilding solicitations analyzed, 17 used a Trade Off approach and 13 used the Lowest Priced, Technically Acceptable (LPTA) approach. Of the procurements that used the Trade Off approach, the technical factors were weighted more important than price (Matsuda et al., 2025).

Of the one Egyptian solicitation for shipbuilding that was analyzed, the sealed-bidding solicitation for four submarines used a point-based evaluation system basis for award (Matsuda et al., 2025).

Of the three JMSDF solicitations analyzed, two solicitations used the proposal-based competition method (similar to FAR-based Trade Off) and one solicitation used the open solicitation method (similar to FAR-based LPTA; Matsuda et al., 2025).

Of the four Taiwan Navy solicitations analyzed, all four used a Most Advantageous Tender approach, similar to the Trade Off approach discussed in the U.S. FAR. In addition, it should be noted that Taiwan Navy solicitations also include the budget amount of the procurement within the solicitation (Yang, 2023).

Source Selection Team Composition

In the U.S. solicitations analyzed, although the source selection team was not specified in the solicitation, the source selection team composition reflected the policy in the DoD FAR Supplement (DFARS) concerning source selection procedures. This included having a formal Source Selection Authority (SSA), Advisory Council (SSAC), and Evaluation Board (SSEB). It should be noted that FAR policy states that participating as a voting member on any source selection board is considered an inherently government function and must be performed by federal government employees. There are no other limitations on number of SST members or the positions (Matsuda et al., 2025).

Although the Egyptian solicitations did not include the source selection team member names, the source selection structure as reflected in the EMOD policy consist of three entities involved in the source selection: the Technical Specifications Committee, Market Research Committee, and the Technical Evaluation Committee (Matsuda et al., 2025).

The Japan solicitation documents did not reveal any information about the source selection team composition (names, number of evaluators). The Japan Acquisition, Technology, and Logistics Agency (ATLA) procurement regulations clearly define the source selection team structure to include the head of the Designated Sole-Source Contract Review Committee (DCSRC) as well as its members. Additionally, the names of high-ranking positions, as defined by the regulations, are published as a list of key executives. Therefore, the names of the head of the evaluation team, the DSCRC chair, and some DSCRC members are available to all interested parties, including offerors. It should also be noted that in one of the Japan solicitations, it stated that fairness and transparency were ensured by having a non-government third party audit the evaluation process. It is unclear whether this non-government third-party requirement existed in the other Japan solicitations as well (Matsuda et al., 2025).

In the Taiwan solicitations analyzed, all of the names of the source selection team members (Procurement Evaluation Committee and Working Group) were published on the agency procurement website designated by the responsible authority. In addition, the Taiwan procurement code specifies a minimum requirement of the total number of source selection team members and minimum percentage requirement for source selection team members outside the government agency (e.g., representatives from industry and academia; Yang, 2023).

Proposal Evaluation Factors

In the U.S. solicitations analyzed, the 13 solicitations using the LPTA approach established proposal evaluation factors with price as the highest priority, and technical acceptability and past performance acceptability as the next priority. For the 17 solicitations that used the Trade Off source selection approach, the evaluation factors reflected technical approach and design as the most important factor, then the management factor, and then the past performance factor and then the price factor, which is the least important. In addition, for the Trade Off approach, the non-price factors, when combined, were considered significantly more important than price (Matsuda et al., 2025).

As previously stated, the Egyptian solicitation used a point-based evaluation system method for source selection. The solicitation included evaluation factors consisting of price factors weighing more important than non-price factors. The price factors, in order of importance, were total cost of acquisition, life cycle cost, and payment structure. The non-price factors, in order of importance, were technical specifications, operational support and maintenance, and technology transfer and local content (Matsuda et al., 2025).

In Japan's solicitations, proposal evaluation criteria are divided into two main categories: mandatory items and additional items. The mandatory items, which are not identified in the solicitation but are based on case data, are related to prospective contractors' ability to perform the contract to stated specifications. The additional items are made up of factors that evaluate more advanced expertise, design, technology, creativity, etc. The additional items are assessed only for proposals that have satisfied all of the mandatory items, and points are awarded according to the proposal evaluation criteria and scoring system that was determined prior to the issuance of the solicitation. The analysis of Japan solicitations covered three solicitations. One of the solicitations (DD) contained no specific descriptions of the proposal evaluation criteria in the RFP. This is no surprise since the RFP stated that the guidelines for preparing the necessary documents to be submitted by offerors will be handed over to them in person. Therefore, there is a possibility that information on the proposal evaluation criteria is included in the guidelines for preparing the necessary documents to be distributed to offerors. Thus, these guidelines were not made public. The solicitation for the AOE clearly stated the proposal evaluation criteria. The criteria included non-price items (shipbuilding facilities, shipbuilding techniques, quality control system, cost reduction measures, training of shipbuilding personnel, status of technological partnerships, status of participation in competitive shipbuilding contracts for other ships, information security capability, shipbuilding past performance, and response to technological challenges related to shipbuilding) and price. The solicitation for the PV included the following evaluation items: advanced naval vessel design and construction, integrated management capability of related companies for onboard equipment, and integrated management capability from design to sustainment and maintenance (Matsuda et al., 2025).

In the Taiwan solicitations analyzed, all four of the solicitations used a ranking method to select the most advantageous tender. The ranking method evaluates and awards a score against each evaluation item and the total scores of all tenders are converted to a ranking. The rankings are totaled to determine each tender's overall ranking. The tender that attains the lowest overall ranking is ranked first. Thus, the ranking method is converted through the score and then totaled with the ranking numbers, and the offeror with the lowest ranking number is the most advantageous tender for the buyer. In addition, Taiwan Navy policy states that if the price factor is included in the consideration of the evaluation, the weighting of the price factor shall be not less than 20% and not more than 50% compared to all of the evaluation items. Thus, the score of the price as an evaluation factor should be no less than 200 and not more than 500 if the total score for all the evaluation factors is 1000. The evaluation factors used were consistent in the four solicitations. The importance of the evaluation factors ranged in descending order from technical capability, price, offeror's organizational structure and financial condition, management and execution of project plan, past performance, other factors related to the functions or benefits of this procurement to performance of presentation and on-site questions and answers (Yang, 2023).

Conclusions and Areas for Further Research

Although the U.S. FAR best value continuum reflects at least three source selection basis for award approaches, LPTA, Trade Off, and Highest Technically Rated Offer (HTRO), in the solicitations analyzed in this research, only LPTA and Trade Off were used. Additionally, the findings indicate that Egypt utilizes the point-based evaluation system (closely resembles Trade Off). Our findings also indicate that for Japan's solicitations, they use two types of source selection evaluation approaches. Proposal-based competition (closely resembles Trade Off), and the open solicitation method (closely resembles LPTA). Unlike Egypt and the United States, Japan does not disclose mandatory proposal evaluation criteria values in its RFPs. It is interesting to note that the Taiwan Navy solicitations included the budget amount of the procurement within the solicitation. This may have implications for ensuring integrity, accountability, and transparency in the source selection process.

In our analysis on Source Selection Team (SST) composition, information about the composition of SSTs was unavailable in all cases, likely because none of these countries require public disclosure of such information, including the names and expertise of evaluators. As a result, our comparative analysis of SST composition in the United States, Egypt, Japan, and Taiwan is based on the regulatory frameworks. In terms of regulations, the first notable difference lies in the composition and regulation of evaluation teams. In the United States, the SSEB includes advisors, cost or pricing experts, legal counsel, small business specialists, and subject-matter experts. Egypt follows similar policies, prohibiting all but government employees from voting membership in SSEBs. In Japan, evaluators may not always operate independently for each evaluation criterion, and voting members are permitted to serve on multiple boards simultaneously depending on the complexity of the contract. In addition, in Taiwan source selections, evaluators may be supplemented by private industry, academia, or any other sector where government employees are lacking sufficient expertise or credentialing.

For the 30 U.S. solicitations, we found that there was consistency in the proposal evaluation criteria within LPTA source selections and Trade Off source selections. For LPTA solicitations, price was the most important evaluation criteria followed by technical specifications and past performance. For Trade Off solicitations, non-price factors such as technical merit or approach and past performance consistently were more important than price. In the Egyptian solicitations, price was the most important criteria followed by non-price factors of technical specifications, operational support and maintenance, and technology transfer. For the Japan solicitations, we could not access the mandatory evaluation items which are not stated in the solicitation, but did identify the additional items to include facilities, techniques, quality control and cost reductions, training, proposed partnerships, information security, and past performance. We did not have any findings about whether Japan incorporates relative importance for either mandatory or additional items. The Taiwan solicitations all four of the solicitations used a ranking method to select the Most Advantageous Tender (similar to Trade Off).

Based on the findings of this research, the following areas for further research are identified for U.S. shipbuilding solicitations. First, the United States should further explore the benefit of using the HTRO source selection approach for our most technically advanced ships (aircraft carriers, destroyers, submarines). Second, the United States should explore if including the procurement budget within the solicitation would increase the accountability, integrity, and transparency of the shipbuilding procurement. Third, the United States should explore if there is any benefit in making the names of the members of the source selection team publicly accessible for interested parties. This approach may also add to the accountability, integrity, and transparency of our shipbuilding procurements. Finally, the United States should further explore the benefits of including nongovernment personnel, such as industry and academia representatives who have needed expertise, as member of the source selection team, in either a voting or non-voting capacity.

References

DoD OIG. (2025). *Summary report: Lessons learned from DoD OIG reports on contract oversight* (Report No. DODIG-2025-096). https://www.dodig.mil/reports.html/Article/4192287/summary-report-lessons-learned-from-dod-oig-reports-on-contract-oversight-repor/

Egyptian Ministry of Defense. (2010). *The official home page of the Egyptian Armed Forces.* https://www.mod.gov.eg/modwebsite/Default.aspx

GAO. (2024). *Navy shipbuilding: Increased use of leading design practices could improve timeliness of deliveries* (GAO-24-105503). https://www.gao.gov/products/gao-24-105503

GAO. (2025a). *Navy shipbuilding: Enduring challenges call for systemic change* (GAO-25-108225). https://files.gao.gov/reports/GAO-25-108225/index.html

GAO. (2025b). *Shipbuilding and repair: Navy needs a strategic approach for private sector industrial base investments* (GAO-25-106286). https://files.gao.gov/reports/GAO-25-106286/index.html

Global Taiwan Institute. (2022, November 30). *Taiwan's naval shipbuilding programs point towards an evolving direction in defense policy*.

https://globaltaiwan.org/2022/11/taiwans-naval-shipbuilding-programs-point-towards-an-evolving-direction-in-defense-policy/

JMOD. (2024). *Reiwa roku nen ban nippon no bouei* [Defense of Japan 2024]. Nikkei insatsu.

O'Rourke, R. (2024). Great power competition: Implications for defense-issues for Congress. https://apps.dtic.mil/sti/html/trecms/AD1219945/

Matsuda, S., Peters, A., & Aboutaleb, A. (2025). *Comparison of source selection approach between U.S., Egyptian, and Japanese shipbuilding procurement.* Naval Postgraduate School, Acquisition Research Program.

Ministry of National Defense, Taiwan. (2021). *National defense report*. https://www.mnd.gov.tw/PublishForReport.aspx?&title=%E8%BB%8D%E4%BA%8B%E5%88%8A%E7%89%A9&SelectStyle=%E6%AD%B7%E5%B9%B4%E5%9C%8B%E9%98%B2%E5%A0%B1%E5%91%8A%E6%9B%B8%E5%B0%88%E5%8D%80

Minnick, W. (2016, June 23). *Taiwan moves on $14.7B indigenous shipbuilding, upgrade projects*. Defense News. https://www.defensenews.com/naval/2016/06/23/taiwan-moves-on-14-7b-indigenous-shipbuilding-upgrade-projects/

National Contract Management Association. (2022). *Contract management body of knowledge* (7th ed.).

Rendon, R.G., & Rendon, J. M. (2015) Auditability in public procurement: An analysis of internal controls and fraud vulnerability. *International Journal of Procurement Management*, *8*(6), 710–730.

U.S. GSA, DoD, & NASA. (2026). *Federal Acquisition Regulation* (48 C.F.R. ch. 1). https://www.acquisition.gov

Yang, S. Y. (2023). *Comparison of source selection strategies between the United States' and Taiwan's shipbuilding procurement*. Naval Postgraduate School, Acquisition Research Program.

Neglect With the Old, In With the New: How Neglect of Existing Weapon Systems Is Causing a Readiness Issue within the Department of War, And How Acquisition Approaches Can be Improved to Solve It

Moshe Schwartz—is President of Etherton & Associates and holds positions as Senior Fellow for Acquisition Policy at NDIA and Director at the Procurement Round Table. He previously served as the Executive Director of the 809 Panel on acquisition reform.

Arzoo Dawoodani—is a second-year student and research assistant at the George Washington University Law School and interning at Etherton & Associates.

Anthony Di Stasio—is a former Senior Executive with the Pentagon where he served as the Director for the MCEIP organization. He was responsible for the running of the DPA Title III and IBAS investment portfolios. He has 20 years of experience with Acquisitions and investments across multiple disciplines including weapons, mineral, micro-electronics, shipbuilding, space and workforce development. Most recently he served as the acting Deputy Assistant Secretary of Defense for Industrial Base Resilience before departing Government service in November 2025.

Abstract

This research identifies the challenges driving poor weapon system operational readiness within the Department of War (DoW) and proposes ways to improve associated readiness rates. Weapon system operational readiness has been on a steady decline over the last two decades. This decline in readiness rates has been highlighted by both the General Accountability Office, the DoW Inspector General, and DoW officials at various levels.

Introduction

The Department of War (DoW; DoD)[1] has a readiness problem. As a retired master sergeant recently quipped, "Maintenance has been in a bad way for 20 years" (Campbell, 2025).

Consider these findings from the Government Accountability Office (GAO):

- From 2011 through 2021, "sustainment challenges worsened" for 10 Navy ship classes, leading to a decrease in the number of ship availability hours for operations or training" (GAO, 2023).
- From 2022 to 2024, Air Force aircraft mission-capable rates declined 4% (Tirpak, 2025).
- In 2024, the Department did not meet "its mission capable rate goals for 42 of the 45 DoD aircraft that support military-related missions" (GAO, 2025, p.17).

The DoD Inspector General agrees. In its report on the seven "top DoD management and performance challenges" for FY26, "Maintaining and Improving Material Readiness" for weapon systems ranked third (U.S. Department of Defense Office of Inspector General [DoD OIG], 2025, p.2). According to the IG, sustainment challenges that impact mission-ready weapon systems affect all the services, spanning aircraft, ships, and vehicles (DoD OIG, 2025, p.3).

[1] Congress, the Inspector General, GAO, and other entities use the term Department of Defense. This paper will use the acronym DoD when referring to agencies that use it (such as the DoD IG), when quoting statute, and when quoting a report that uses the acronym DoD.

For its part, the Department recognizes the problem. In his Questions for the Record, Air Force Chief of Staff Kenneth Wilsbach wrote “our ability to generate forces has declined significantly in the past five years” (Senate Armed Services Committee [SASC], 2025, p.10).

What is Driving Poor Readiness Rates

According to the GAO, Navy officials identified lack of spare parts, lack of trained maintenance personnel, and increases in deferred maintenance as the primary drivers of the readiness challenge. These long-running problems were found to prevent maintenance during deployment, increase the cost of repairs, and reduce ship service life and operational readiness (GAO, 2025b). In other reports, the GAO found that maintenance challenges directly impacting Navy readiness include aging equipment, cancelled/deferred maintenance (due to not fully funding maintenance accounts), reliability of ship systems, and inefficient shipyard layout (an infrastructure issue; GAO, 2023b, 2024a, 2024b; SASC, 2025).

The DoD IG identified three focus areas to sustain “mission-ready capabilities and ensure rapid responsiveness in any contested environment”

- Maintenance,
- Storage and Upkeep, and
- Logistics Networks (DoD OIG, 2025; SASC, 2025).

Aging equipment and inadequate maintenance were also identified by the DoD IG as DoD-wide challenges (DoD OIG, 2025; Mitchell Institute, 2026, 21:10). In some instances, the sustainment strategy has been identified as a barrier to better readiness rates (Bergman, 2026).[2]

Delayed depot-level maintenance and modernization are the most significant limiters of current Fleet readiness

CNO Fighting Instructions, 2025, p. 16

The DoW recognizes these challenges to readiness. For example, in its *Fighting Instructions: 2025*, the Chief Naval Officer identified multiple sources of readiness challenges, including delayed/deferred maintenance, drydock space, maintenance plans that are not optimized, insufficient work package planning, and workforce needs (U.S. Navy, 2026, p. 16).

Such challenges lead to increased maintenance demands, cannibalizations of parts (removing working parts from one platform and using them elsewhere due to parts shortages), or maintenance delays. As Table 1 indicates, cannibalizations and maintenance delays (days beyond the scheduled end date for depot maintenance) have each increased, while steaming hours has decreased (GAO, 2023a).

[2] “The V-22 Osprey faces readiness challenges due to its complex design, high maintenance manhour requirements, inefficient supply system, and current sustainment strategy. These factors have led to persistently low mission-capable rates and a 30% increase in operating and maintenance costs per flight hour over the past four years” (Bergman, 2026).

Table 1. Changes in Sustainment Metrics per Ship Across Selected Navy Ship Classes, Fiscal Years 2011 through 2021

Ship class	Maintenance cannibalizations[a]	Days of maintenance delay
Ticonderoga-class cruiser (CG-47)	+3 ▲	+7 ▲
Nimitz-class aircraft carrier (CVN-68)	+4 ▲	+7 ▲
Arleigh Burke-class destroyer (DDG-51)	+7 ▲	+20 ▲
Freedom-class littoral combat ship (LCS-1)	+15 ▲	0 ●
Independence-class littoral combat ship (LCS-2)	+3 ▲	+19 ▲
America-class amphibious assault ship (LHA-6)[b]	-1 ▼	0 ●
Wasp-class amphibious assault ship (LHD-1)	+9 ▲	+10 ▲
San Antonio-class amphibious transport dock (LPD-17)	+3 ▲	+33 ▲
Whidbey Island-class dock landing ship (LSD-41)	+6 ▲	+19 ▲
Harpers Ferry-class dock landing ship (LSD-49)	+7 ▲	-16 ▼
Fleetwide	**+6** ▲	**+14** ▲

● No change (neutral) ▲ Increase (negative) ▼ Decrease (positive)

Source: U.S. Government Accountability Office, "Weapon System Sustainment: Navy Ship Usage Has Decreased as Challenges and Costs Have Increased." GAO-23-106440 Report to Congressional Committees (January 2023). https://www.gao.gov/assets/gao-23-106440.pdf

According to the Navy's deep dive into the V-22 program, missed readiness targets are due to four main factors:

1) not sharing and implementing known maintenance best practices;
2) poor supply systems and maintenance programs that don't prioritize program readiness outcomes;
3) reliability issues, outdated publications, and not complying with published standards; and
4) inventory management challenges, including delayed deliveries (NAVAIRSYSCOM, 2025, p. 12).

Interestingly, in the reports, audits, instructions, opening statements, and questions for the record discussed above, intellectual property or technical data rights are not raised as a significant readiness or maintenance challenge.

Drilling Down into Readiness Drivers

Workforce

If the DoW had all the spare parts are on hand, maintenance facilities that were fully modernized, perfect visibility into maintenance requirements, supply chains that worked perfectly, and all the IP, the Department would still face readiness challenges. Without sufficient numbers of skilled personnel to conduct maintenance, sustainment does not get done.

Without sufficient numbers of skilled personnel, sustainment does not get done.

DoW personnel, including maintainers, are overworked and understaffed (Campbell, 2025; GAO, 2024c; Mitchell Institute, 2026, 23:00). The Navy has stated that it lacks sufficient trained maintenance personnel to maintain combat surface ship (GAO, 2025b, p.9). Half of the Air Force bomber units have personnel shortages for key roles, including maintenance (SASC, 2025, p. 41). This issue was highlighted by John Venable, CEO and Resident fellow of the Mitchell Institute for Aerospace Studies, when he stated, "I don't see us being able to fix the airplanes unless we have the number of crews that can actually be qualified and be available to fix them" (Mitchell Institute, 2026b, 23:42).

Remote operations can alleviate workforce concerns by allowing for maintenance and installation of equipment of government personnel qualified by contractors, assisted remotely by contractors. Such an approach could expedite maintenance in the field, without having to bring contractors in from distant locations or shipping equipment back to contractor facilities. Using specific contractor–qualified personnel could also keep warranties in place. Remote-assisted operations could also reduce the Department's own workforce manpower requirements by leveraging a similar approach to maintenance and installation.

Underfunded Maintenance and Infrastructure Accounts

Another challenge is not fully funding maintenance, spare parts, or related infrastructure accounts, which leads to maintenance backlogs, deferred maintenance, lack of parts, and insufficient preventative maintenance.

Underfunded Infrastructure

An aging maintenance infrastructure, legacy equipment, and outdated processes lead to insufficient capability to keep up with maintenance requirements (Mohan, 2026). The organic industrial base performs a substantial amount of remanufacturing and maintenance; the average organic industrial base factory is 80 years old (Temin, 2024). Stephanie Hoaglin, Director of the Army's Organic Base Modernization Task Force stated that the Army has an $18 billion, 15-year plan to upgrade infrastructure, improve equipment and process, and "bring all those things to the 21st century" (Temin, 2024). This means that if all goes according to plan, the Army organic industrial base will achieve 21st century status some 40 years after the century started.

The organic industrial base can significantly increase readiness with facility upgrades and improved use of technology. While some upgrades, like the Army's full 15-year plan, require substantial funding, in the short term, certain less-expensive upgrades could prioritized to appreciably increase maintenance capacity.

The Navy has similar challenges, reporting that without improvements to shipyard infrastructure, it will be unable to support almost a third of the planned maintenance periods for aircraft carriers and submarines through 2040. According to the GAO, dry docks and maintenance facilities are poor, and equipment is "generally past its useful life" (GAO, 2023b, p.1). Shipyard infrastructure in the domestic commercial does not have the capacity to support national security needs. Fewer than 1% of commercial ships are built domestically. The United States has only 66 shipyards (including 8 shipbuilding yards and 22 repairs yards with drydocking; The White House, 2026).

The Trump administration promulgated Executive Order 13807 Establishing Discipline and Accountability in the Environmental Review and Permitting Process for Infrastructure to expedite the permitting process for infrastructure investment) and Executive Order 14269 Restoring America's Maritime Dominance (to revitalize shipbuilding), but these efforts require significant resources and commitment to succeed.

Deferred/Underfunded Maintenance

Maintenance accounts that are not fully funded lead to deferred maintenance and insufficient preventative maintenance. In 2023, the Navy's surface ship maintenance backlog was $1.8 billion (GAO, 2023c, p. 19). Included in the underfunded accounts are those dedicated to buying spare parts. For example, the Air Force unfunded spare parts requirement was recently estimated to be $8.9 billion (SASC, 2025, p.43). In a fireside chat at the Mitchell Institute, Lt. Gen. David A. Harris, Deputy Chief of Staff, Air Force Futures stated that if he had extra dollars, “the first dollar I would spend would be on spare parts . . . your deployed spares packages” (Mitchell Institute, 2026a, 21:20).

Spare Parts

Even when the trained personnel are available, the existing infrastructure supports maintenance needs, and funding is available to buy spare parts, too often the spare parts are not on the shelf. As Lt. Gen. Kenyon Bell, Air Force Deputy Chief of Staff for logistics, recently stated

> *When maintainers are available, they are trained, they are ready, and then they reach back and there is not anything on the shelf in order to replace a broken asset, that is a problem. (Mitchell Institute, 2026b, 22:28)*

AF Chief of Staff Wilsbach also identified acquiring sufficient spare parts for weapon systems as a challenge, adding that having a “healthier inventory of parts” could improve weapon system availability (SASC, 2025, p. 10, 42).

Three drivers of spare parts shortages that keep coming up in our research are

- insufficient communication with industry/the organic industrial base,
- poor data management/IT systems, and
- outdated acquisition processes.

Communication

The DoW has a communication problem. Demand signal challenges and the breakdown of communication between the Pentagon and industry was the subject of a paper coauthored by Moshe Schwartz, *What We’ve Got Here Is Failure to Communicate: How Better Communication Can Improve DoD Acquisition Outcomes* (Schwartz et al., 2025). Secretary of War Hegseth echoed a similar theme (perhaps he read our paper?), when commenting on defense acquisition he stated:

> *We have to fix our own house first—provide clarity, simplify the system, allow more people to access it, give that steady demand signal* (emphasis added). (Zeljko, 2026)

The Air Force is working to improve communication and more effectively sending a spare parts demand signal to the industrial base. According to Lt. Gen. Bell, “We are making sure that the both the organic industrial base and the defense industrial base understand the demand signal for each of the components and the parts that are needed . . . so that maintainers have what they need” (Mitchell Institute, 2026b, 22:43).

Effective IT Systems and Data Management

Poor data management and outdated IT systems have a surprising impact on readiness. One challenge is the sheer number of duplicative, outdated systems that often do not talk to each other. According to Army undersecretary Mike Obadal, the Army has 42 different systems for training and readiness, which means that “commanders don’t have a cohesive picture of readiness, and to get close requires hundreds of man hours to compile data” (Miller, 2026). The Army also has some 75 different logistics systems, adding to the weapon system readiness challenge (Miller, 2026). Not all of these systems are effectively interoperable.

In one example, the Army relies on the Defense Logistics Agency (DLA) for 90% of its maintenance parts—mostly low-dollar, high-volume expendables (Williams, 2025). Yet until recently, the Army and DLA operated with incompatible systems, resulting in supply chain disruptions and hindering readiness. DLA was unable to assess what supplies the Army was consuming and unable to order the necessary parts before they were needed, a significant challenge when some parts have a two-year delivery time (Williams, 2025).

The current demand planning accuracy rate
"doesn't help the readiness of our systems"

Maj. Gen. David Sanford

Another challenge is the accuracy and effectiveness of IT systems, in part due to their age. The DLA has challenges with its demand planning accuracy, which is reportedly at approximately 60% (Obis, 2026). DLA director of logistics operations, Maj. Gen. David Sanford, acknowledged that the current demand planning accuracy rate "doesn't help the readiness of our systems" (Obis, 2026).

In another example, DoD IG's analysis of spare parts on Navy ships found that the spare parts "inventory accuracy was between 83% and 95%, which is below the minimum inventory accuracy of 98% needed to ensure the ships' readiness" (DoD OIG, 2025b, p.3). The inventory discrepancies in the Navy's RSupply software application (intended to provide real-time, online tools for inventory management) were attributed to personnel not knowing where the parts were, not updating inventory records after issuing spare parts, and not removing excess line items from the ships. As a result, the IG found that "the Navy did not have assurance that the 10 ships we reviewed in the Indo-Pacific region had all of the required spare parts . . . to maintain operational readiness" (DoD OIG, 2025b, p.3).

These issues are exemplified in the most recent Inspector General's Audit of the Army's Management of Repairs. This audit found that the Army did not "adequately maintain BFV's [Bradley Fighting Machines], turned in non-mission capable BFV's, and improperly reported the condition of the BFVs turned in as fully mission capable" (DoD OIG, 2026, p.6). This ineffective management was also the result of the unit's lack of personnel and capabilities to repair the BFV's. This is a lack of communication, lack of spare parts, and ineffective IT systems at play.

Outdated Acquisition Processes

The Department could increase the availability of spare parts by improving acquisitions, including increasing reliance on commercial parts and using multiyear procurement processes.

Buy Commercial

Too often, the Department acquires parts and components that are bespoke and customized when commercial products could meet requirements. Reevaluating specific requirements could lead to identifying parts that have unnecessary requirements, are over-designed, or obsolete and replaceable with alternatives. Special Operations Command has been successful in looking for commercial solutions to incorporate into the design process, thereby avoiding customized parts or components.

Use Multiyear Procurements

Sometimes, the Department has used a just-in-time approach of ordering parts, placing orders only when there are specifically identifiable maintenance needs. This approach engenders unnecessary delays, as maintainers await the arrival of parts or components.

Multiyear procurements can ensure regular delivery of parts in predetermined quantities, ensuring that the parts are always available.

Multiyear procurements offer other benefits, including driving down cost through larger quantity buys, giving industry a stronger demand signal that could incentivize contractors to invest in infrastructure, and create a more stable production rate.

The Department in recent months has been pursuing multiyear procurements to improve acquisition and readiness issues.

Use all the Tools in the Toolkit

The Department has various tools at its disposal to incentivize industry and outside capital. Some administrations have focused on direct investment and grants, to varying degrees. Some have focused on fixed-price contracts, others on cost-contracts. This administration is leaning heavily on loan guarantees, equity investments, and insisting on industry investment in facilities. Focusing on value to the government instead of profit margin of the supplier is another valuable tool that is not leveraged enough, particularly when industry has invested significant resources to develop new solutions.

Too often, the Department took a single approach and applied to almost all situations—as if a silver bullet exists. There is no silver bullet, no single approach to incentivize industry. The DoW must leverage all the tools available to it, customized to the particular situation instead of pursuing one-size-fits-all approach.

Aging Weapon Systems

Aging equipment is yet another challenge. Older equipment needs more maintenance, increasing cost, adding downtime, and reducing readiness. The average aircraft in the Air Force is approximately 30 years old, compared to an average age of 17 years in 1994 (SASC, 2025). Age of systems also contributes to spare parts shortages. Some of the weapon systems are so old that the companies that made certain parts are no longer in business and left no technical data packages behind. This is why earlier this year, the Department announced a $1.8 million investment to develop reverse engineering capabilities to develop tech data packages for these parts (DoW, 2026). While this investment is a positive step, more needs to be done.

Are Intellectual Property and Data Rights Major Drivers of Poor Readiness Rates

Some have argued that the DoD's limited data rights to the intellectual property needed for sustainment is the primary barrier to improved readiness. The data tell a different story.

While inadequate data rights can sometimes be a factor, an analysis of GAO reports, DoD IG audits, independent analyses, and DoW statements cited in this paper indicate that intellectual property is not the driver of the readiness crisis.[1] Intellectual Property (IP) is not even raised as an issue in most of these sources.

The DoD has the legal authority and updated policy guidance (including the creation of the IP cadre) necessary to effectively negotiate for, manage, and use data rights and IP. The DoW's problems are rooted in how it manages IP. As the Joint Explanatory Statement for the FY2026 National Defense Authorization Act Stated

> *The Department's challenges related to technical data are not rooted in an insufficiency in the law, but rather insufficiencies in the Department's planning and resourcing decisions made early in the acquisition phase related to the sustainment of the systems it procures, and in some cases the Department's insufficient inspection, acceptance, and management of technical data that have been negotiated**

The DoW also lacks a systematic method to track IP rights across the services and even between programs, leading to wasteful and duplicative licensing of the same technology by different programs

A Path to Improving Readiness

Many of the issues affecting readiness are systemic. They are also fixable. The DoW and Congress have already taken steps within the last two years to tackle some of these issues in a serious way. We modestly add a few recommendations below as the Department continues down the path to improving readiness.

Plan Sustainment Early

Too often, the Department does not plan for sustainment early in the acquisition process. Better and earlier planning will help

- design systems for faster or more cost-effective sustainment,
- identify early-on the Department's IP needs. Such early identification will allow for better contracts (see Intellectual Property section below) and let industry know what IP might be required by the Government, allowing industry to build that into their business case.

1. *Fully implement of section 1803 of the FY26 NDAA*

The issue of insufficient early planning for sustainment and bringing the maintainers in early in the design/acquisition process was partially addressed in section 1803 of the FY26 NDAA (Life Cycle Management and Product Support) which elevated the product support manager to be coequal with the program manager and requires the product support manager to plan for IP requirements and management. The provision also requires

- product support managers meet certain certification and training requirements,
- sustainment reviews to focus on why systems are not meeting operational requirements and readiness objectives,
- sustainment plans to include strategies for public-private capabilities, plans for IP management, and using best-value approaches in life cycle planning, and
- the Office of Cost Assessment and Program Evaluation to maintain a database on O&S estimates and actual costs for major weapon systems

Effective implementation of section 1803 could significantly improve readiness. That is why we strongly encourage the Department and Congress to do so.

Leverage IT to Improve Sustainment and Maintenance Management

From upgrading legacy systems to predictive analytics, the DoW and Congress could take steps to leverage technology to improve readiness.

2. *Implement predictive analytics capabilities to identify what maintenance can be conducted and parts ordered <u>before systems break down</u>*

The GAO has highlighted the need for the DoD to further implement predictive maintenance, which could increase operational availability of weapon systems (GAO, 2022). Such an approach can include embedding artificial intelligence within weapons systems to flag maintenance problems (U.S. Navy, 2026, p. 14). While the Department has begun to pilot these predictive maintenance programs within weapon systems, the GAO has found that there is no action for acting on forecasts being provided (GAO, 2022, p.2; Thompson, n.d.).

The Department recently introduced a new metric called materiel resilience to capture how systems are performing under stress and how quickly it can return to baseline (Miles et al., 2026, p. 37). While this could give clear insight into sustainment issues early on and support proposed predictive analytics efforts, this data would need to be implemented across all platforms and monitored regularly to ensure a proactive approach to system sustainment.

3. *Establish a MITIER (Modernizing IT Infrastructure to Expedite Readiness) tiger team to accelerate modernizing the IT infrastructure to manage spare parts and maintenance sequencing*

Be it spare parts management on surface ships, spare parts forecasting at the DLA, or the Army's multitude of logistics and readiness systems, the IT infrastructure responsible for managing spare parts and maintenance is in desperate need of an upgrade. Yes, we are talking about business systems, but these are business systems that **directly correlate to mission capable readiness rates of weapon systems**. Modernizing IT systems to ensure data reliability, interoperability, and forecasting will help ensure that the Department has the right spare parts, delivered to the right locations, at the right time. This is not a just-in-time approach but a must-be-in-time concept to ensure parts are where they need to be **before they are needed.**

We recommend the Department, through the office of the Chief Data and AI Officer, in coordination with the DLA and the military services, establish a tiger team specifically charged with modernizing those systems responsible for managing the spare parts (including setting a demand signal to industry) and maximizing the efficiency of maintenance sequencing.

4. *Leverage technology to improve sustainment during operations*

The Department is leveraging technology to streamline and improve

- operational planning and deployment,
- forecasting for intelligence and predictive analytics for maintenance, and
- logistics.

To date, these efforts appear to be pursued individually. The Department could integrate these tasks into an integrated capability that enables operational planning informed by

- AI-enabled predictive analytics estimating sustainment requirements from battle damage and operational tempo *based on actual data and planned actions,* and
- current logistics capabilities and global stock/maintenance facility availability.

Integrating these tasks into a single capability could also speed up the process for activating logistics and sustainment needs based on real-time operational planning and deployments data.

Improve IP Management Within the Department

Even if IP and Data Rights are not significant drivers of poor weapon system readiness, the Department has IP challenges that should be addressed. Secretary Hegseth recently said about acquisition, "We have to fix our own house first—provide clarity, simplify the system, allow more people to access it." These three principles—clarity, simplifying the system, and more shared access—are a framework for improving IP in the Department.

5. *Include detailed, specific IP asset schedules in contracts similar to those used in commercial agreements.*

The DoW does not always include detailed and specific IP asset schedules into contracts. As a result, industry and government do not have a true meeting of the minds. This causes unnecessary confusion and disagreements later on.

6. *Implement Section 805 of the FY2026 NDAA*

The Department lacks a centralized record keeping system for IP licenses. Section 805 of the FY26 NDAA requires the Department to implement a system to track covered data and contractor compliance with requirements for technical data. Implementing this system will help the Department better understand and manage its IP and give the DoW the visibility it needs to ensure that contractors are held accountable for complying with IP agreements.

Section 805 also required the Senior Acquisition Executives to review contracts to identify where they didn't receive, can't find, or didn't order needed IP, and enter into active negotiations to get those insufficiencies in legacy systems resolved. If implemented correctly, this will also help the Department manage IP and ensure that it contracts for the IP it needs.

7. *Protect Against At-Risk Subcontractors in the Supply Chain*

Prime contractors, and by extension the DoW, rely on numerous sole-source manufacturers that own the IP for critical parts or services necessary to support weapon systems acquisition and sustainment.

Some of these manufacturers are not financially healthy and at risk of being acquired by adverse capital, going out of business, or being acquired by a company that will simply increase the cost to DoW. Sometimes, the company may be healthy, but the particular product line is not profitable, prompting the subcontractor to phase out or divest/sell the product line. In these scenarios, the continued supply is prohibitively expensive and often not an attractive business proposition for industry.

One option is to create a mechanism for a non-government holding entity to

- acquire financially distressed manufacturers producing essential parts, components, or systems, and
- acquire the IP, tech data, and other information necessary to manufacture parts, components, or systems that are being phased out in industry and not supported by the manufacturer.

For at-risk ,sole-source manufacturers supporting multiple programs and subcontracting to multiple primes, the entity could acquire the company and provide the parts to the primes as Government Furnished Equipment (GFE).

This entity could collaborate with the Office of Strategic Capital (OSC), in partnership with OUSW(A&S) Industrial Base Policy, DIU, DCMA, DLA, PAEs, sustainment centers, depots, etc. to identify and prioritize buying a critical supplier or acquiring the IP needed to self-manufacture when necessary. In many cases, prime contractors are best positioned to identify at-risk critical suppliers. A revolving fund can be used to fund the initial investment and, when necessary, with future revenue streams used to replenish the revolving fund.

To succeed, the Department may need to pursue thoughtful intellectual property and technical data polies for legacy parts and small manufacturers. One approach could be to include a right of first refusal or a grant/purchase of certain IP rights for government purposes in the event of a sale of the company, or the phasing out of supporting an item in industry. Limiting the use of acquired IP to defense applications would protect commercial markets while securing national interests.

Invest in the Workforce and Contractor-DoW Teamwork

As discussed above, there are insufficient numbers of maintainers and other sustainment critical personnel. Retention remains a challenge (GAO, 2026). In some cases, there is insufficient training for maintenance personnel.

8. *Invest the resources to build a sufficiently sized, trained, and experienced workforce*

Understanding the costs involved, we believe the investment in a sufficiently sized and skilled workforce is justified by an anticipated higher rate of weapon system readiness. Such a workforce cannot be turned on quickly, and investing in a sustainment workforce today will be a critical node in ensuring that in a war with a peer adversary tomorrow, the United States has the capacity to keep our Air Force climbing high, Navy anchors away, and Army rolling along.[3]

9. *Implement experimenting with remote contractor maintenance assistance to reduce costs, downtime, and improve contractor-warfighter teamwork*

In certain circumstances, remote contractor maintenance assistance to enhance maintenance and readiness. Such an approach may only be applicable in limited circumstances—but could still be a useful approach. In such an approach, contractors could “certify” certain DoW personnel to conduct maintenance *when done in coordination and with remote-participation of contractors.* This approach could reduce wait times for qualified contractor personnel to arrive onsite and work on systems, decrease travel costs, and improve training capabilities for onsite personnel.

[3] Musical reference intended.

Appendix I: Recommendations

Source	Title	Recommendations
Department of Defense Office of the Inspector General	2026 Report	1. Retire non-compliant financial systems, which would resolve audit discrepancies and save about $760 million annually (DoD OIG, 2026, p. 39)
		2. Increase the number of trained maintenance personnel (DoD OIG, 2026, p.13).
		3. Implement proactive management to increase oversight of storage and upkeep (DoD OIG, 2026, p. 18)
Department of the Navy and NAVAIRSYSCOM	V-22 Osprey Hearing	1. Establish reliability control boards to address readiness degraders (NAVAIRSYSCOM, 2025, p. 16)
		2. Implement a supply cell to address interoperability and consolidate parts allocation across services (NAVAIRSYSCOM, 2025, p. 15)
		3. Implement closed loop dealing for maintainers to leverage their experience and keep them on for longer periods of time (NAVAIRSYSCOM, 2025, p. 14)
		4. Conduct cross-service supply diagnostics to identify and tag any related issues across systems (NAVAIRSYSCOM, 2025, p. 15)
Mitchell Institute	AirPower Forum	1. Partner with industry and integrate unique capabilities (Mitchell Institute, 2026a, 17:17)
		2. Invest enough to maintain surge capacity for munitions and platforms for wartime (Mitchell Institute, 2026a, 39:01)
		3. Implement the Aircraft Readiness Machine framework for interoperability (Mitchell Institute, 2026b, 16:14)
		4. Bring maintainers in but keep them at the generalist area for a bit to improve their skill set and quality of work (Mitchell Institute, 2026b, 26:47)

Senate Armed Services Committee	Advance Policy Questions for General Kenneth S. Wilsbach, USAF	1. Shift responsibility to senior leadership to pay attention to maintenance of weapon systems
		2. Invest in spare parts availability so that the shelf is full and there are less maintenance delays
U.S. Navy	Chief Naval Fighting Operations Instructions	1. Imbed artificial intelligence into core functions to increase operational capability (U.S. Navy, 2026, p. 14)
		2. Fix backlog in munitions production by implementing multi-year procurement authorities, produce with allies, optimize production priority across customers, and partner with industry (U.S. Navy, 2026, p. 15)
Government Accountability Office GAO-23-10556	Military Readiness: Actions Needed to Further Implement Predictive Maintenance on Weapon Systems	1. All services should develop a comprehensive implementation plan for predictive maintenance. This includes action plans, milestones, outcome-related goals, and a framework to keep track of everything. (GAO, 2022)
Government Accountability Office GAO-23-106440	Weapon System Sustainment: Navy Ship Usage Has Decreased as Challenges and Costs have Increased	1. Implement a more comprehensive and mission-specific readiness tracking. DoD should broaden its targets for when ships are available, and broaden what counts as taking a ship out of commission to include unplanned maintenance, unplanned losses, and training (GAO, 2023b).
Government Accountability Office GAO-23-106673	Military Readiness: Improvement in Some Areas, but Sustain and Other Challenges Persist	1. DoD should prioritize establishing metrics for measuring readiness to conduct full-spectrum operations across domains (GAO, 2023c, p. 4).
		2. DoD should update its F-35 sustainment strategy for the supply chain (GAO, 2023c, p. 15)
		3. Navy should focus on obtaining accurate cost estimates to optimize facilities, replace aged equipment, and improve overall Shipyard infrastructure (GAO, 2023c, p. 27)
Government Accountability Office GAO-25-106728	Amphibious Warfare Fleet: Navy Needs to Complete Key Efforts	1. Secretary of the Navy should establish a time frame for their ongoing joint plan to address ship

	to Better Ensure Ships are Available for Marines	availability concerns (GAO, 2024b) 2. Navy amphibious ship depot maintenance policy should be updated to include that depot maintenance should not be cancelled before amphibious ships have reached the end of their expected service life (GAO, 2024b) 3. Chief of Naval operations should establish performance goals with tangible objectives to measure progress (GAO, 2024b)
Government Accountability Office GAO-24-105917	Military Readiness: Comprehensive Approach Needed to Address Service Member Fatigue and Manage Related Efforts	1. Assessment of DoD's oversight structure for fatigue related efforts. Assessment should identify and delegate authority to an office to act as a focal point for and oversee DoD fatigue-related efforts (GAO, 2024c, p.39) 2. Senior leadership should create and maintain a comprehensive list of all fatigue-related research projects (GAO, 2024c, p.39)
Government Accountability Office GAO-25-106990	Navy Surface Ships: Maintenance Funds and Actions Needed to Address Ongoing Challenges	1. Set availability requirements that capture all factors that could contribute to ships being needed before they're ready (GAO, 2025b, p.10). 2. Regularly report on assessments of risks associated with deferred maintenance on surface ships (GAO, 2025b, p.10).
Government Accountability Office GAO-25-108104	Military Readiness: Implementing GAO's Recommendations Can Help DoD Address Persistent Challenges across Air, Sea, Ground, and Space Domains	1. Analyze data to identify high-risk training activities and ensure guidelines reflect conditions affecting the needed personnel and training (GAO, 2025a, p.9). 2. "Determine desired mix of government and contractor roles" (GAO, 2025a, p.24). 3. Create ship industrial base strategy, "prioritize strategy to not prematurely cancel maintenance when divesting ships" (GAO, 2025a, p.35). 4. "Complete planning elements before fielding new equipment" (GAO, 2025a, p.35).

References

Bergman. J. (2026, February 10). *Joint SPF/RDY hearing: V-22 Osprey program update*. House Armed Services Committee. https://armedservices.house.gov/calendar/eventsingle.aspx?EventID=6392

Campbell, A. (2025, July 25) *Decades of troubles for Air Force maintainers set to get worse with job consolidation*. Military.com. https://www.military.com/daily-news/investigations-and-features/2025/07/25/decades-of-troubles-air-force-maintainers-set-get-worse-job-consolidation.html

DoD, OIG (2025b, May 14). *Evaluation of the spare parts onboard U.S. Navy ships in the Indo-Pacific region* (Report No. DODIG-2025-100). https://media.defense.gov/2025/May/15/2003715601/-1/-1/1/DODIG-2025-100_REDACTED_FINAL_SECURE.PDF

DoD, OIG. (2025a, November 25). *Top DoD management and performance challenges fiscal year 2026*. https://media.defense.gov/2025/Nov/25/2003831751/-1/-1/1/MANAGEMENT%20CHALLENGES%20FY2026.PDF

DoD, OIG. (2026, March 16). *Audit of the Army's management of repairs to Bradley Fighting Vehicles to meet U.S. Army Europe and Africa Mission requirements* (Report No. DODIG-2026-065). https://media.defense.gov/2026/Mar/18/2003900233/-1/-1/1/(U)%20DODIG-2026-065%20FINAL%20REPORT%20REDACTED.PDF

DoW. (2026, February 9). *DOW addresses material obsolescence through reverse engineering training*. https://www.war.gov/News/Releases/Release/Article/4401952/dow-addresses-material-obsolescence-through-reverse-engineering-training/

GAO. (2022, December 8). *Military readiness: Actions needed to further implement predictive maintenance on weapon systems* (Report No. GAO-23-105556). https://www.gao.gov/products/gao-23-105556

GAO. (2023a, January 31). *Weapon system sustainment: Navy ship usage has decreased as challenges and costs have increased* (Report No. GAO-23-106440). https://www.gao.gov/products/gao-23-106440

GAO. (2023b, June 28). *Navy readiness: Actions needed to address cost and schedule estimates for shipyard improvement* (Report No. GAO-23-106067). https://www.gao.gov/products/gao-23-106067

GAO. (2023c, May 2). *Military readiness: Improvement in some areas, but sustainment and other challenges persist* (Report No. GAO-23-106673). https://www.gao.gov/assets/gao-23-106673.pdf

GAO. (2024a, May 1). *Military readiness: Actions needed for DoD to address challenges across the air, sea, ground, and space domains* (Report No. GAO-24-107463). https://www.gao.gov/assets/gao-24-107463.pdf.

GAO. (2024b, December 3). *Amphibious warfare fleet: Navy needs to complete key efforts to better ensure ships are available for marines* (Report No. GAO-25-106728). https://files.gao.gov/reports/GAO-25-106728/index.html

GAO. (2024c, March 26). *Military readiness: Comprehensive approach needed to address service member fatigue and manage related efforts* (Report No. GAO-24-105917). https://www.gao.gov/assets/gao-24-105917.pdf

GAO. (2025a, March 12). *Military readiness: Implementing GAO's recommendations can help DoD address persistent challenges across air, sea, ground, and space domains* (Report No. GAO-25-108104) https://www.gao.gov/assets/gao-25-108104.pdf

GAO. (2025b, January 31). *Navy surface ships: Maintenance funds and actions needed to address ongoing challenges* (Report No. GAO-25-106990). https://www.gao.gov/products/gao-25-106990

GAO. (2026, March 4). *Military readiness: DoD should take further actions to address challenges across the air, sea, ground, and space domains* (Report No. GAO-26-108888). https://www.gao.gov/products/gao-26-108888

Miles, J., Cruz, F., Estelle, C. Broadwell, G., Gulick, K., & Weinstein, D. (2026, February). Lessons from a joint exercise at the point of need. *Defense Acquisition Magazine.* *https://online.flippingbook.com/view/1055717789/4/*

Miller, J. (2026, February 23). Army tackling its "Achilles heel" of IT modernization. *Federal News Network*. https://federalnewsnetwork.com/ask-the-cio/2026/02/army-tackling-its-achilles-heel-of-it-modernization/

Mitchell Institute. (2026a, February 19). *Fireside chat with Lt. Gen. David A. Harris, Deputy Chief of Staff, Air Force Futures*. YouTube. https://www.youtube.com/watch?v=soF0CTOe1Rc

Mitchell Institute. (2026b, January 29). *Air Force readiness*. YouTube. https://www.youtube.com/watch?v=xuup_PUCdYY

Mohan, C. (2026, February 24). *Joint TAL/RDY hearing: Modernization of the organic industrial base*. House Armed Services Committee. https://armedservices.house.gov/calendar/eventsingle.aspx?EventID=6397

NAVAIRSYSCOM. (2025, December 12). *V-22 comprehensive review*. Department of the Navy. https://www.secnav.navy.mil/foia/readingroom/HotTopics/V-22%20Review/V-22%20Comprehensive%20Review%20(Distro%20A).pdf

Obis, A. (2026, January 21). DLA turns to AI, ML to improve military supply forecasting. *Federal News Network*. https://federalnewsnetwork.com/defense-main/2026/01/dla-turns-to-ai-ml-to-improve-military-supply-forecasting/

Schwartz, M., Johnson, M., & Schwartz D. (2025, May 5). What we've got here is failure to communicate: How better communication can improve DoD acquisition outcomes (SYM-AM-25-309). In *Proceedings of the Twenty-Second Annual Acquisition Research Symposium and Innovation Summit, Volume I.* https://dair.nps.edu/bitstream/123456789/5358/4/SYM-AM-25-309.pdf

Temin, T. (2024, October 17). How the Army goes about modernizing its crucial but aging organic industrial base. *Federal News Network*. https://federalnewsnetwork.com/army/2024/10/how-the-army-goes-about-modernizing-its-crucial-but-aging-organic-industrial-base/

Thompson, N. In pursuit of failure: US military and predictive maintenance. *Global Defence Technology*. https://defence.nridigital.com/global_defence_technology_apr24/in_pursuit_of_failure#nav-area.

Tirpak, J. (2025, February 18). Air Force mission capability rates reach lowest levels in years. *Air and Space Forces Magazine*. https://www.airandspaceforces.com/air-force-mission-capable-rates-fiscal-2024/

U.S. House of Representatives, Committee on Armed Services. (2025) *Joint explanatory statement to accompany the National Defense Authorization Act for Fiscal Year 2026.* *https://armedservices.house.gov/uploadedfiles/fy26_ndaa_joint_explanatory_statement.pdf*

U.S. Navy. (2026, February 9). *Chief naval operations fighting instructions, 2025*. https://media.defense.gov/2026/Feb/06/2003871752/-1/-1/1/CNO%20FIGHTING%20INSTRUCTIONS.PDF/CNO%20FIGHTING%20INSTRUCTIONS.PDF

U.S. Senate, Committee on Armed Services. (2025, October 9). *Advanced policy questions for General Kenneth S. Wilsbach, USASF nominee for appointment to be Chief of Staff of the Air Force*. https://www.armed-services.senate.gov/imo/media/doc/wilsbach_apq_responses.pdf

The White House. (2026, February 13). *Restoring America's maritime dominance.* https://www.whitehouse.gov/maritimemight/

Williams, L. (2025, January 17). Is bad data to blame for missing weapons parts? *Defense One*. https://www.defenseone.com/defense-systems/2025/01/bad-data-blame-missing-weapons-parts/402324/

Zeljko, R. (2026, February 10). "Impossible to deal with": Pete Hegseth reveals the real culprit behind defense contractor delays. *Blaze Media.* https://www.theblaze.com/news/impossible-to-deal-with-pete-hegseth-reveals-the-real-culprit-behind-defense-contractor-delays

PANEL 21. OPPORTUNITIES TO RAPIDLY DELIVER WARFIGHTING CAPABILITIES USING LEADING PRACTICES

Thursday, May 7, 2026, 1445 – 1600 ET (1145 - 1300 PT)

Panel Summary:

In an era of rapid technological change, the Department of Defense (DOD) must move beyond linear acquisition models to maintain a decisive edge. This panel explores the critical shift toward iterative development and agile portfolio management, highlighting how leading commercial companies prioritize investments to deliver value at speed. By examining Government Accountability Office (GAO) assessments of weapon system testing and technology investment oversight, these researchers provide an informative baseline for operationalizing reforms that transform the test and evaluation enterprise from a reactive bottleneck into a proactive catalyst for modernization.

Chair: Shelby Oakley, Director, Contracting & National Security Acquisitions, U.S. Government Accountability Office

Panel Presenters:

Weapon Systems Testing: DOD Needs to Update Policies to Better Support Modernization Efforts – *Christopher Durbin, Assistant Director, U.S. Government Accountability Office (GAO)*

Leading Practices: Agile Portfolio Management and Iterative Business Cases Drive Innovative Product Development – *Brenna Derritt, Assistant Director, Contracting and National Security Acquisitions, U.S. Government Accountability Office (GAO)*

Defense Research and Engineering: Action Needed to Improve Management and Oversight of Technology Investments – *Brian Smith, Senior Analyst, U.S. Government Accountability Office (GAO)*

Shelby Oakley—is a Director in the Government Accountability Office's (GAO) Contracting and National Security Acquisitions team. In her role, she oversees GAO's portfolio of work examining the most complex and expensive acquisitions within the federal government. Her portfolio includes Navy Shipbuilding and Nuclear Triad recapitalization programs, DOD acquisition policy and oversight, and leading practices in product development. In addition, she is responsible for GAO's annual work to assess the cost, schedule, and performance of DOD's portfolio of major defense and middle-tier acquisition programs.

Further, her portfolio also includes oversight of acquisition management at the Veterans Affairs Department. Ms. Oakley previously served as a Director in GAO's Natural Resources and Environment team where she led teams reviewing a range of nuclear security, policy, and nonproliferation related issues.

From 2004 to 2015, Ms. Oakley led teams reviewing the activities of the National Aeronautics and Space Administration (NASA) with a focus on helping NASA improve its acquisition management practices. Her reviews covered key aspects of NASA's operations, such as Space Shuttle workforce transition and sustainment of the International Space Station, as well as reviews of all major NASA systems including in-depth reviews of NASA's human spaceflight programs and the James Webb Space Telescope.

Ms. Oakley earned a Master's Degree in Public Administration from the University of Pittsburgh's Graduate School of Public and International Affairs and her Bachelor of Arts Degree from Washington and Jefferson College.

Weapon Systems Testing: DoD Needs to Update Policies to Better Support Modernization Efforts

Shelby S. Oakley—is a director in the Government Accountability Office's (GAO) Contracting and National Security Acquisitions team. In her role, she oversees GAO's portfolio of work examining the most complex and expensive acquisitions within the federal government. Her portfolio includes Navy Shipbuilding and Nuclear Triad recapitalization programs, DOD acquisition policy and oversight, and leading practices in product development. In addition, she is responsible for GAO's annual work to assess the cost, schedule, and performance of DOD's portfolio of major defense and middle-tier acquisition programs. [oakleys@gao.gov]

What the GAO Found

The Department of Defense (DoD) identified test and evaluation modernization as a crucial part of its effort to get capabilities to warfighters faster. DoD organizations, including the Office of the Secretary of Defense and the military departments, have undertaken modernization planning efforts with varying areas of focus and levels of detail. Nonetheless, these plans share themes, including the use of digital engineering tools and highly skilled workforces.

The GAO's analysis of **DoD-wide test and evaluation policies** found they were not fully consistent with selected leading practices for product development as applied to test and evaluation: involve testers early, conduct iterative testing, use digital twins and threads, and obtain user feedback iteratively. These policies contained some tenets of the leading practices, particularly for the software acquisition and urgent capability acquisition pathways. However, these leading practices were largely not reflected in the policies for programs in the major capability acquisition and middle tier of acquisition pathways, which account for the majority of DoD spending on weapon systems acquisition.

Key DoD-wide Weapon Systems Testing Policies Fall Short of Selected Leading Practices

Leading practices for product development applied to test and evaluation	Department of Defense (DOD) test and evaluation policy: Current developmental test policy		Current operational and live-fire test policy	
	Overall	MCA and MTA	Overall	MCA and MTA
Involve test officials in development of acquisition strategy	◐	MCA ○ MTA ◐	◐	◐
Conduct iterative testing to generate minimum viable product	◐	○	◐	○
Use digital twins and digital threads	◐	○	◐	◐
Obtain user feedback throughout iterative testing	◐	○	◐	○

◐ Policy partially reflects leading practice ○ Policy does not reflect leading practice

MCA = Major capability acquisition
MTA = Middle tier of acquisition
} Majority of DOD spending on weapon systems acquisition

Overall = overall assessment of test policy for four weapon system acquisition pathways: MCA, MTA, urgent capability acquisition, and software acquisition

Source: GAO analysis of DOD-wide policies applicable to test and evaluation. | GAO-26-107009

Further, the GAO found that the **DoD's digital engineering policy** and the test and evaluation section of the **DoD's systems engineering policy** do not describe specific processes to ensure application of leading practices to testing.

The GAO also found that **military department-level test and evaluation policies** generally did not reflect the leading practices beyond the level found in DoD-wide policies. The GAO similarly found that these leading practices were not reflected in key program documents, like acquisition strategies and test strategies, for selected weapon systems acquisition programs it reviewed.

The DoD has a unique opportunity to not only retool its existing test and evaluation enterprise, but to redefine the role that enterprise can play in enabling faster delivery of relevant capabilities to warfighters. Fully incorporating leading practices into policies relevant to weapon system test and evaluation could help pivot the test enterprise's current reactive role to a proactive one, informing and aiding defense acquisition efforts.

Why the GAO Did This Study

The DoD has yet to realize its goal to rapidly develop weapon systems to get capabilities to the warfighter when needed. DoD acquisition programs have identified challenges discovered during test and evaluation as contributing to delays in development.

A House committee report includes a provision for the GAO to assess how the DoD is modernizing weapon system test and evaluation. The GAO's report (1) describes the DoD's plans to modernize test and evaluation to deliver capabilities faster to the warfighter, and (2) assesses the extent to which DoD-wide and military department policies for test and evaluation reflect selected GAO leading practices for product development.

To do this work, the GAO assessed DoD test and evaluation modernization plans and policies and weapon system acquisition documentation. The GAO visited three military department test organizations to observe tools in practice. The GAO also interviewed DoD and military department officials from test organizations and other entities.

What the GAO Recommends

The GAO is making 13 recommendations, including that the Secretary of Defense and the Secretaries of the Air Force, Army, and Navy each should revise weapon system test and evaluation policies and other related policies to reflect selected leading practices for product development. Specifically, revisions should require involvement of testers in acquisition strategies; iterative approaches to testing, including use of digital twins and threads; and ongoing end user input. The DoD concurred with seven recommendations, partially concurred with five recommendations, and did not concur with one recommendation. The GAO continues to believe all 13 of its recommendations are valid, as discussed in this report.

Note: this is a highlights-level summary of a full GAO report. See GAO-26-107009 for additional details.

Leading Practices: Agile Portfolio Management and Iterative Business Cases Drive Innovative Product Development

Brenna Derritt—is an Assistant Director in the GAO's Contracting and National Security Acquisitions team, with more than 10 years of experience assessing Department of Defense (DoD) acquisition programs and contracts. She is currently co-leading the GAO's annual assessment of weapon system acquisitions. Previously, she led and supported multiple reviews that identified commercial practices for product development that could benefit government acquisitions. Brenna received a Master of Public Policy degree from the University of Chicago and a Bachelor of Arts degree from the University of Colorado Boulder. [derrittb@gao.gov]

Abstract

This report (GAO-25-107130) is the third of a series on product development leading practices that are instructive for improvements to the Department of Defense's acquisition processes. The Government Accountability Office's (GAO) recent work has emphasized the importance of structuring defense acquisition programs around iterative development--a technical process crucial to how leading companies develop innovative, value-added products on timelines responsive to users' needs. GAO's work has found that leading companies employ equally robust business processes to the management of individual product developments as well as the overall product mix, or portfolio. Understanding how leading companies decide on their product development investments and how they use business cases--justifications for undertaking a product development--to support their decisions can inform ongoing acquisition reform efforts. GAO identified eight leading companies based on rankings in well-recognized lists, interviewed, company representatives, and analyzed documentation.

Background

For more than 25 years, the U.S. Government Accountability Office (GAO) has recommended actions that the Department of Defense and other federal agencies should take to improve their most complex and costly acquisition programs, saving taxpayers tens of billions of dollars. Most recently, this work has emphasized structuring acquisition programs around iterative development—a process stemming from Agile software development that leading companies use to develop innovative, value-added products that respond quickly to users' needs. Such leading practices offer proven approaches that can inform improvements to agencies' acquisition of complex systems.

A sound technical process is not the only driver for the consistent pattern of success that characterizes leading companies' product development. Leading companies employ equally robust business processes to initiate, justify, and prioritize their investments in innovative products such as semiconductors and industrial automation systems.

Leading Companies Use Agile Portfolio Management Practices to Guide Product Development Investments

Leading companies employ a forward-looking, agile approach to managing the overall mix of products, or product portfolios, that enables their iterative product development processes. Through agile portfolio management, leading companies continuously evaluate product development opportunities and investments using three key recurring practices (see below).

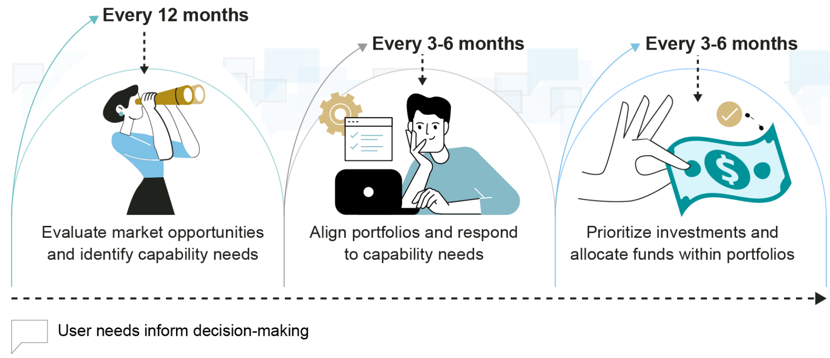

Source: GAO analysis of leading company information; PureSolution/stock.adobe.com. | GAO-25-107130

In addition, portfolio managers continually interact with business cases—collections of information that justify undertaking product development efforts. For example, based on regular assessments of business case data, portfolio managers may decide to add resources to improve weaker performing products or discontinue outdated products that impede demand for updated versions. Continually updating portfolios based on business case data enables leading companies to optimize their investments and ensure portfolios are responsive to the company's strategic vision and evolving user needs.

Leading Companies Use Iterative Business Cases to Inform Scaled Investment in a Product's Development

Leading companies recognize that traditional business processes are ill-suited to support the rapid, iterative technical approaches that guide modern product development. Consequently, these companies have implemented iterative business processes, including for product business cases themselves. Leading companies now ensure that initial business cases for new products address four key elements—market and user needs, product definition, internal value, and target schedule. These companies then update business case data throughout product development to inform recurring decisions on product priorities (see below).

Initial business case	Identifies the market and user needs the product will fulfill.	Defines the product, including the key capabilities it will provide to users.	Assesses the net, internal value the product will provide to the company.	Outlines a target schedule for delivering a minimum viable product (MVP).
Updated iterations of business case	Monitors evolving market and user needs and assesses whether the planned product is situated to fulfill those needs.	Reviews technology developments and innovations and assesses whether the product's definition remains relevant to users.	Revisits and refines estimates of the net, internal value the product will provide to the company.	Revisits and refines the target schedule for delivering an MVP.

Source: GAO analysis of leading company information; PureSolutions/stock.adobe.com. | GAO-25-107130

Importantly, leading companies scale their investment in a product's development based on the pace at which that product's business case matures, rather than on promised performance or the passage of time (see below).

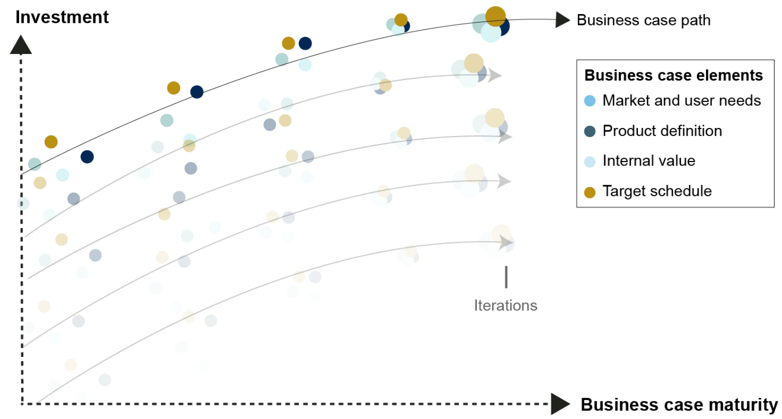

Source: GAO analysis of leading company information. | GAO-25-107130

This is an excerpt from the full-length report. See GAO-25-107130 for details, including additional report contributors: https://www.gao.gov/products/gao-25-107130 (report landing page); https://www.gao.gov/assets/gao-25-107130.pdf (PDF of report).

Defense Research and Engineering: Action Needed to Improve Management and Oversight of Technology Investments

Brian Smith—is a Senior Analyst with the U.S. Government Accountability Office's Contracting and National Security Acquisitions team. For more than 10 years in his 19-year career at the Government Accountability Office, he has worked on research and engineering issues at DOD, reporting and making recommendations to improve how DOD manages its technology development efforts. [smithbt@gao.gov]

Abstract

The Department of Defense (DOD) seeks to outpace foreign adversaries' capabilities by quickly adopting innovative technologies. The Office of the Under Secretary of Defense for Research and Engineering (OUSD[R&E]) has responsibility for managing, overseeing, and improving technology development efforts across DOD to help reach that goal. In the President's fiscal year 2026 budget submission, DOD requested nearly $180 billion for research, development, test and evaluation (RDT&E) activities aimed, in part, at developing technologies that meet both the short-term and long-term needs of current and future warfighters. This request included more than $20 billion for science and technology (S&T) activities and more than $40 billion for advanced component development and prototyping efforts, funding for which OUSD(R&E) is responsible for providing management and oversight.

OUSD(R&E) and the Office of the Under Secretary of Defense for Acquisition and Sustainment were established in February 2018 following the dissolution of DOD's Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics (OUSD[AT&L]). OUSD(R&E)'s duties and responsibilities include serving as DOD's Chief Technology Officer; establishing policies on, and supervising, all defense research and engineering, technology development, technology transition, prototyping, and experimentation activities and programs; and designating senior officials for critical technology areas supportive of the National Defense Strategy, among others.

OUSD(R&E) Is Generally Implementing Programs and Processes Consistent with Its Authorities to Manage and Oversee Technology Investments

In response to its statutory and policy authorities, OUSD(R&E) enacted programs and processes to manage and oversee technology investments. For example, the office developed and released the 2023 National Defense Science and Technology Strategy aligned with the 2022 National Defense Strategy, as required by statute. The strategy is anchored by 3 strategic pillars—mission focus, foundation building, and succeeding through teamwork—which were then translated into three strategic lines of effort to establish the ways to sharpen DOD's competitive edge: (1) focus on the Joint Mission; (2) create and field capabilities at speed and scale; and (3) ensure the foundation for research and development. The military departments developed their own science and technology strategies, but these strategies do not fully align with DOD's strategy. There is no requirement, in policy or statute, for the military departments to update their science and technology strategies or to align their strategies to the National Defense S&T Strategy. We found several areas where the military departments' strategies do not align with DOD's overarching strategy. Updating and aligning the military departments' S&T strategies to the maximum extent practicable with the National Defense S&T Strategy would allow the military departments and OUSD(R&E) to ensure a common vision for technology development across DOD.

In addition, consistent with its authorities, OUSD(R&E) also designated senior officials who oversee critical technology areas (CTA), and it initiated processes for conducting technology reviews and collecting technology transition data. OUSD(R&E) is also administering several prototyping programs, meant to quickly deliver technologies to the warfighter.

OUSD(R&E) Faces Challenges Managing and Overseeing Military Department Technology Efforts

OUSD(R&E) faces challenges in managing and overseeing military department technology development efforts. For example, it has yet to ensure that CTA roadmaps consistently provide sufficient information for military departments to invest in technologies for the joint fight. The roadmaps vary, in part, because OUSD(R&E) has not developed formal guidance for the CTA Principal Directors—the senior officials designated to coordinate research and engineering activities and develop research and technology development roadmaps for each CTA—to use in developing the roadmaps. By not issuing guidance for roadmap development, OUSD(R&E) lacks reasonable assurances that the military departments are focusing on and investing in technologies considered critical to meeting the NDS and maintaining technological superiority against its adversaries. In addition, OUSD(R&E) officials acknowledged they have not provided guidance for roadmap development to the military departments. For example, this would include not identifying the level of military department investments that OUSD(R&E) would consider necessary to ensure alignment to the maximum extent practicable with each CTA. Without directing the level of investment needed in each CTA, OUSD(R&E) lacks reasonable assurance that sufficient investments are being made towards progress in any given CTA to ensure timely delivery of future capabilities. Further, OUSD(R&E) risks insufficient investments being made in the technologies it has identified as being critical to countering the threats of our adversaries.

OUSD(R&E) also lacks statutory authority to confirm that the military departments' technology investments, as expressed in their annual budget submissions, align with OUSD(R&E) priorities, and is limited in its ability to influence military department RDT&E budgets to ensure they align with department-wide priorities. Without proper authority and time to review and assess RDT&E budget submissions as part of the budget process, DOD risks the military departments not investing in technologies that warfighters need both for the current and future fight, especially to support the joint force. In addition, having a complete understanding of the full breadth of technology efforts undertaken by the military departments will enable OUSD(R&E) to provide effective management and oversight of these efforts as well as enable it to provide information to Congress as part of the budgeting process.

This is an excerpt from a full-length report. See GAO-26-107664 for additional details, including additional report contributors: https://www.gao.gov/assets/gao-26-107664.pdf

PANEL 22. DIGITAL FORTIFICATION: MASTERING SOFTWARE ASSURANCE AND CYBER SURVIVABILITY

Thursday, May 7, 2026, 1445 – 1600 ET (1145 - 1300 PT)

Panel Summary:

In today's contested environments, the reliability of warfighting systems is determined by the quality and assurance of the code that powers them. This panel explores the critical shift toward Software-Defined Warfare, addressing the urgent need for robust Software Management Plans (SMPs) and data-driven cyber survivability requirements. By examining a new acquisition process for military simulations and a repeatable model for decomposing complex cyber attributes into verifiable engineering specifications, these researchers provide the framework necessary to ensure U.S. forces field systems that are not only capable but dependable and resilient against evolving digital threats.

Chair: Gaurang Dävé, Cyber Technology Officer, Marine Corps System Command (MCSC)

Panel Presenters:

> **Stop Chasing the Perfect Requirement Specification: Formalizing Conceptual Model Documentation in Simulation Acquisition –** *Lieutenant Colonel Matthew Morse, PhD, United States Marine Corps*
>
> **Operationalizing Cyber Survivability Through Requirements Decomposition: A Marine Corps Case Study –** *Kathleen Coen, Computer Scientist, Marine Corps Tactical Systems Support Activity*
>
> **Acquisition Software Management Planning (SMP) is Critical to Mission Success –** *Carol Woody, Principal Researcher, Software Engineering Institute*

Gaurang Dävé— serves as Cyber Technology Officer for the Marine Corps Systems Command (MCSC). In this role, he serves as the principal cyber advisor and technical expert to the Commander, Executive Director and Chief Engineer, MCSC; and Program Executive Office Land Systems (PEO-LS) and subordinate activities, for Cyberspace, Cybersecurity, and information technology (IT).

Mr. Dävé previously served as Cyber Advisor at MCSC. In this role, Mr. Dävé provided technical cyber Subject Matter Expertise (SME) for cyber policy, budgeting, evaluating new technologies, representing command senior leadership at DoD, DoN, and within USMC for cyber related matters.

Prior to joining MCSC, Mr. Dävé served as Senior Cyber Technical Advisor and Cybersecurity Program Director for Naval Surface Warfare Center (NSWC) Dahlgren. In this capacity, Mr. Dävé led the Navy's combat systems cybersecurity portfolio across the Warfare Systems Program Office. Mr. Dävé provided technical leadership, guidance, and program management leadership. Prior to that, Mr. Dävé served as Cybersecurity Risk Posture lead for the Office of the Chief of Naval Operations, OPNAV N2N6 Cybersecurity Division at the Pentagon. His contributions included synchronizing cyber strategy, standards and requirements, evaluating and prioritizing investments, providing oversight, and resource sponsorship.

Mr. Dävé began his career as a software engineer and served in technical, program management, and supervisory positions across the joint services. Mr. Dävé led the Chemical, Biological, Radiological (CBR) Analysis, Testing, and Systems Engineering Branch at NSWC, Dahlgren Division. Under his leadership, the branch provided SME in CBR modeling & simulation, threat analysis, executed White House initiatives for Medical Countermeasures, and led several ACAT-3 efforts. Mr. Dävé served as CBR Thrust Area Manager for Major Defense Acquisition Program at Joint Science & Technology Office led by US

Army. Prior to joining government service, Mr. Dävé worked in private industry at International Business Machines (IBM). As a Security Software Engineer, Mr. Dävé led several software security program development efforts designed to protect the company's network.

Mr. Dävé holds Bachelor of Science degrees in Computer Science and Biochemical Pharmacology from State University of New York at Buffalo. Master of Science degree in Systems Engineering with focus in Operational Research from George Mason University. Graduate degree in Cyber Engineering from University of Maryland. He is DAWIA Level III certified in Systems Engineering.

Stop Chasing the Perfect Requirement Specification: Formalizing Conceptual Model Documentation in Simulation Acquisition

LtCol Matthew M. Morse, PhD—MOVES Institute, Naval Postgraduate School, United States Marine Corps. [matthew.morse@nps.edu]

Susan K. Aros, PhD—Operations Research Department, Naval Postgraduate School. [skaros@nps.edu]

LTC Charles Rowan, PhD—MOVES Institute, Naval Postgraduate School, United States Army [charles.rowan@nps.edu]

Abstract

Across the U.S. military, modeling and simulation capabilities are increasingly sought to support analysis, experimentation, and training for the employment of complex capabilities in multi-domain operations. Many of the associated simulation acquisition efforts result in failure, despite extensive expenditures of manpower and funds. Efforts to avoid these outcomes have included rigorous specification of requirements and attempts to implement the Software Acquisition Pathway, with limited success. This paper highlights challenges associated with military simulation acquisition and recommends a new approach grounded in an understanding of simulation design and development best practices. To ensure a simulation's requirements and operating context are adequately understood, the simulation designer must provide the requirements owner, and other stakeholders, explicit documentation of the simulation conceptual model. Once the simulation conceptual model has been validated by the requirements owners, it serves as a blueprint for acquisitions partners in the development of the simulation. In addition to providing a clear guide for the development of the simulation, this process will also ensure the delivery of conceptual model documentation which is critical for supporting simulation verification and validation, use, and maintenance. With the Department of War undergoing a massive reevaluation of the acquisitions process, now is the time to revise the acquisition process for simulation design and development.

Keywords: Simulation, Conceptual Model Documentation, Requirements Transition, Acquisition

Introduction

Across the Department of War (DoW), virtual and constructive simulations are increasingly relied upon to support analysis, experimentation, acquisition, testing, training, mission engineering, and digital engineering. These include standalone simulation capabilities as well as distributed simulation environments, ranging from live, virtual, constructive (LVC) environments to simulation-supported wargaming. Increasing complexity of computer simulations and distributed simulation environments compounds challenges for the acquisition of new simulation capabilities, as well as management of existing simulations, increasing risks of failure for simulation acquisition programs.

While all acquisition efforts can face difficulties with the transition from requirements to capabilities, the fundamental nature of simulation design and development presents unique challenges. To ensure suitable and satisfactory design of simulations, it is not enough for requirements documents to specify the scope of the simulation and the required capabilities. Information must also be provided to specify how various entities, entity interactions, and phenomena will be represented. Requirements managers must specify what abstractions or simplifications, and levels of functional/visual fidelity, are necessary or acceptable relative to an intended use. The conventional JCIDS requirements documents (e.g., capability development document) are inadequate for conveying this complex and nuanced information.

It is well established in the modeling and simulation (M&S) domain that conceptual model documentation is a critical tool for ensuring a simulation is designed appropriately (Robinson, 2012). Conceptual model documentation also serves other important roles throughout a simulation's life cycle. This documentation is a prerequisite for execution of verification, validation, and accreditation (VV&A) activities, which are mandated for DoW simulations (DoD, 2024); facilitates more structured management of simulation parametric data and maintenance over the simulation life cycle; and provides the insights needed for users (analysts, wargamers, and training designers) regarding how to leverage a simulation in accordance with its capabilities and limitations (Aros, 2025).

Despite the important role of conceptual model documentation for mitigating risks to simulation acquisition programs, and the explicit requirement of it for simulation VV&A across the DoW, it is rare for a simulation acquisition program to generate conceptual model documentation. Even in the rare instance where conceptual model documentation is specified as a deliverable, limited understanding of conceptual modeling practices by acquisition professionals and industry partners can present obstacles to the delivery and use of conceptual model documentation in the acquisition and VV&A processes.

Legacy Requirements Specification and Transition Process for Simulation

The Joint Capabilities Integration and Development System (JCIDS) (defined by Chairman of the Joint Staff Instruction [CJCSI] 5123.01J) has historically provided the processes through which capability requirements were defined by requirements managers (RMs), and approved by various authorities for transition to acquisition program managers (PMs). The CJCSI specified requirements documents which needed to be produced in support of these processes, including the Initial Capabilities Document (ICD) and Capability Development Document (CDD). The CDD was the primary requirements document which was transitioned to acquisitions professionals (PMs), though it could be supplemented by additional documents.

The introduction of the Software Acquisition Pathway (DoD, 2020) provided an alternative acquisition approach for software under JCIDS and was developed to support software development best practices, replacing the CDD with a Capability Needs Statement (CNS). Unlike the CDD, the CNS is intended to be updated regularly as the software iterates through user testing and feedback, progressing from a minimum viable product to the fielded capability and subsequent refinements.

Under these JCIDS processes, the acquisition PM coordinated with a simulation developer to build an envisioned simulation based on the requirements documents. Unfortunately, these requirements documents were often inadequate for informing the myriad decisions which must be made by simulation developers, and there was no structured way for the requirements owner to ensure that the PM and developer appropriately interpreted the requirements and the intended use once the requirements were transitioned to the acquisition PM. In addition to meeting more tangible requirements for simulation functions, user interfaces, interoperability, and performance capabilities, simulation developers must make many decisions regarding how to abstract their representation of reality for numerous domains (e.g., terrain, electromagnetic spectrum, weather, logistics, human behavior, communications; Abdelmegid et al., 2022; Brooks & Wang, 2015). Decisions must be made about how to achieve different levels of visual and functional fidelity and what assumptions are acceptable relative to the intended use as they design a simplified, altered representation of reality.

The software acquisition pathway can potentially mitigate some of the issues that may arise during the PM / developer interpretation of requirements, since it requires an early minimum viable product to be produced and made available for end user testing and iterative refinement. However, even under this approach there is no requirement for structured communication between the developer and the RM regarding the design of the simulation. Rather, the end users are only able to provide feedback regarding interface design and functions after almost all of the simulation design decisions have been made and, quite literally, codified in the computer simulation. Even domain subject matter experts will be limited in their ability to evaluate and provide feedback regarding most developer decisions in the design of the simulation and representation of entity attributes, behaviors, and interactions.

Ensuring that the simulation design decisions align with the intended use for a simulation requires extensive structured interaction between the RM, users, PM, and the developer. Conducting such reviews requires an intermediate product which can be understood by all parties and which lies somewhere between the requirements specification and the coded simulation. This intermediate product, often referred to as the blueprint for the simulation, is known as the simulation conceptual model (SCM), communicated via the SCM documentation (SCMD). Without documentation of the conceptual model that explains how various entity behaviors, interactions, and phenomena are represented under the hood of a simulation, along with the model logic and algorithms, the RM will be unable to assess the potential for undesirable effects of different design decisions made by the developer. A simulation provided without conceptual model documentation to explain internal design and logic is referred to as a "black box," presenting monumental challenges for management, implementation, VV&A, use, and maintenance of the simulation.

It is important to note that the JCIDS is now defunct, and has been replaced with the Joint Force Requirements Process (JFRP; DoD, 2025b). This change is intended to decrease bureaucracy and increase speed in the acquisition process, by removing unnecessary steps and emphasizing value-added activities early in the acquisition process. Many questions remain regarding the implementation of this transition, particularly for acquisitions processes at the service level and below. The JFRP presents the services with more freedom to refine their acquisitions processes. It remains to be seen how many JCIDS-style processes and artefacts will continue to be maintained by the services. In light of this, this paper continues to refer to JCIDS processes for a discussion of how they may be replaced to support simulation acquisition.

Simulation Conceptual Modeling in the Design and Development of Simulations

An SCM is the abstracted representation of the system or thing to be simulated (i.e., simuland). An SCM "describes what is to be represented, the assumptions limiting those representations, and other capabilities needed to satisfy the user's requirements" (Borah, 2003, p.35). The development and documentation of an SCM is key to the effective and efficient development of a computer simulation model, enabling coordination, collaboration, and identification of gaps prior to the coding of the model (Abdelmegid, et al., 2022; Brooks & Wang, 2015). It is also essential for the verification and validation (V&V) of the resulting simulation model. While it is possible for an SCM to exist merely in the mind of a modeler (Aros, 2025; Robinson et al., 2015), the documentation of an SCM is the tangible form of the SCM that facilitates the communication about, and use of, the SCM. Therefore, it is also important to explicitly discuss the SCM documentation where appropriate.

Development of an SCM is an integral part of the simulation model development process. Figure 1 provides a modeling and simulation development framework (Tolk et al., 2013) positioning the development of the conceptual model just prior to the development of the simulation. Importantly, this figure depicts the loop-back arrows that indicate ongoing V&V checks throughout the M&S development process. Development and refinement of an SCM provides a mechanism for collaboration between stakeholders to determine the capabilities of the envisioned simulation (Robinson, 2008). The documentation of the SCM facilitates this collaboration. The inherent expectation of collaboration in the development of an SCM is clear in a NATO definition of the SCM as “an implementation-independent description of the content and internal representations that represent the sponsor’s, user’s and developer’s combined concept of the system or simulation under development including logic, architecture, algorithms, available data and explicitly recognising [*sic*] assumptions and limitations” (NATO, 2012, p. J-5).

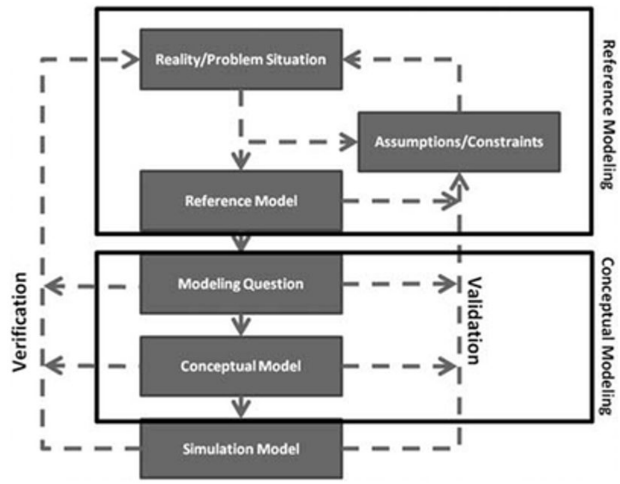

Figure 1. M&S Development Framework (Tolk et al., 2013, p. 6)

The diagram shown in Figure 2 provides additional details on the position of the SCM relative to other artifacts in the simulation development process. Figure 2 also uses dotted-line “comparison” arcs to depict how different artifacts are compared with each other during V&V efforts throughout the development process. Notably, this figure explicitly represents the different roles of requirements and the conceptual model of the simulation.

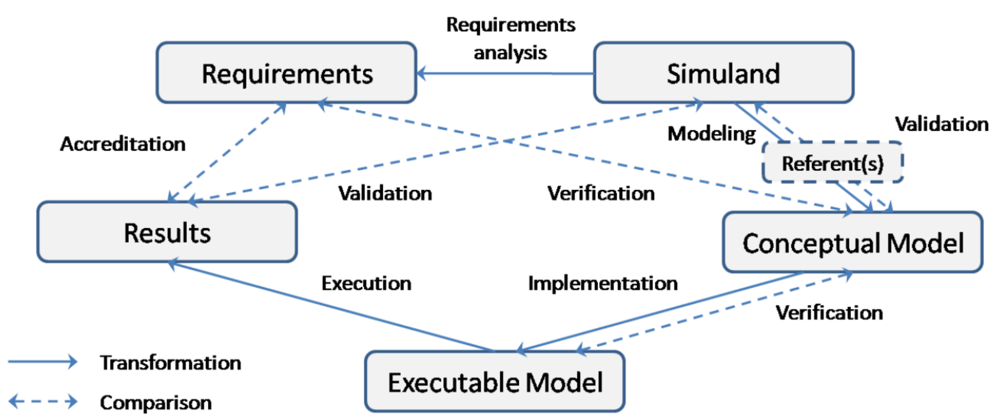

Figure 2. V&V Comparisons Using the Conceptual Model (Appleget, 2011, p. 29, from Petty, 2009)

Requirements Hand-Off during the Acquisition of Military Simulations

Under the JCIDS process, simulation conceptual model development has generally been viewed as occurring after the requirements have transitioned from the RM to the acquisition PM and developers, as seen in Figure 3.

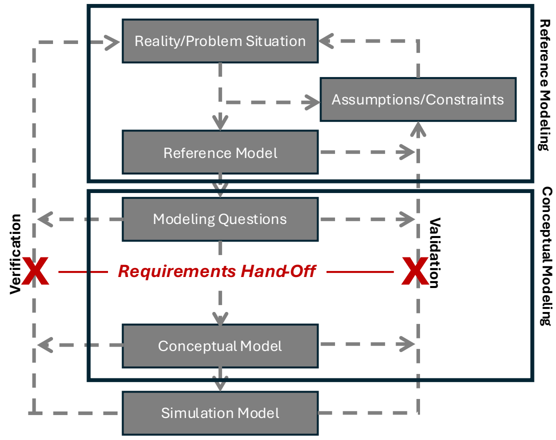

Figure 3. Adapted from M&S Development Framework (Tolk et al., 2013); Shows Effects of a Requirements Hand-Off prior to Conceptual Model Development

This hand-off of simulation requirements, and the lack of an explicit acquisition requirement for SCM documentation, often results in the conceptual model being viewed as a low-priority deliverable rather than a critical tool for informing the design and development of the simulation capability. Consequently, RMs attempt to produce written requirements documents that fully specify all simulation requirements, while also attempting to comply with the "Goldilocks principle" of requirements specification (not too much, not too little). This requirements hand-off also requires the acquisition professional to develop a conceptual model based on their interpretation of the requirements document. This is at odds with established best practices for conceptual modeling in simulation design and development, negatively impacting both acquisition and long-term management of DoW simulations. The Distributed Simulation Engineering and Execution Process (DSEEP) explicitly states that requirements generation should occur after conceptual modeling is conducted: "As the conceptual model is developed, it will lead to the definition of a set of detailed requirements for the simulation environment" (IEEE, 2022).

The steps in the conventional requirements hand-off from the RM to the PM / developers are depicted in Figure 4 using a swim-lane process flowchart to clearly delineate roles and responsibilities. Under the JCIDS process, after requirements documents were generated by the RM and transitioned to the acquisition PM, the RM had no mechanism for ensuring the requirements were correctly interpreted; rather, it was up to the PM / developers to decide whether they would request additional clarification of the requirements from the RM. This can be seen in Figure 4, where the decision diamond appears in the "PM / Developers" swim-lane.

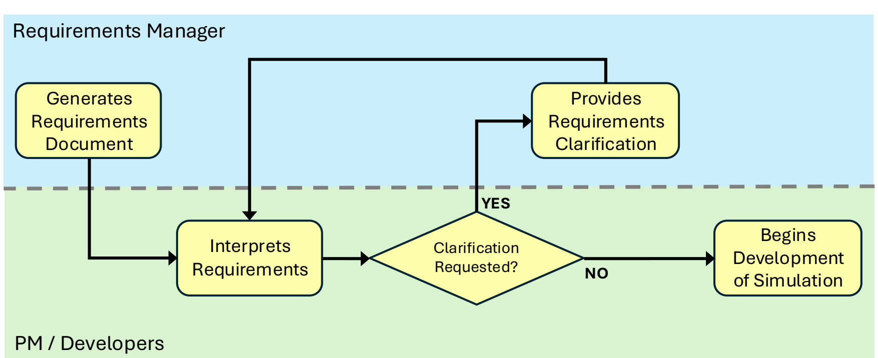

Figure 4. Existing Simulation Requirements Transition Process

For conventional military systems (e.g., vehicles, weapon systems) miscommunication in the requirements transition process can be mitigated to some extent through efforts to enhance clarity and specificity in capability requirements documents. However, for computer simulations, it is more challenging to clearly and concisely specify system capabilities and user interface considerations that include a comprehensive description of the referent, the intended use, and requisite/acceptable abstractions and simplifications. The difficulties associated with succinctly, accurately communicating this information is why conceptual model documentation is so critical for simulation design and development.

Under the JCIDS requirements generation and transition process, RMs are told to avoid providing too much detail in their requirements, which might restrict acquisition professionals' ability to explore the potential solution space. At the same time, RMs aren't provided a feedback loop that would allow them to ensure appropriate interpretation of requirements specification by their acquisition counterparts. This results in an over-specification paradox where RMs are faced with conflicting demands. If they provide too great a level of detail, they restrict acquisitions professionals' abilities to find or craft a desirable solution. If they provide too little, the requirements may be misinterpreted, yielding a capability inappropriate for the intended use. Requirements managers faced with this situation have attempted to find a way to craft simulation requirements that will succeed under these conflicting demands. Meanwhile, the tool that can significantly mitigate the risk of requirements misinterpretation for simulations, SCM development and documentation, has been underemployed across the DoW.

Current State of SCM in DoW Simulation Acquisition

Despite the well-established utility of conceptual model documentation for supporting development, maintenance, and employment of M&S capabilities, it is rare to find SCM documentation for DoW M&S capabilities (Pace, 2011). This lack of SCM documentation presents challenges across several stages in the acquisition life cycle, including maintenance and modernization, simulation database management, and informing user employment of simulations. Conceptual model documentation provides an implementation-agnostic representation of how entities and entity interactions are, or will be, represented by a simulation. They provide insight into how a simulation represents entities, interactions, and phenomena, and specify the logic and algorithms of the simulation. This documentation is a critical resource for supporting simulation design and development of both individual and distributed simulations, and is referenced as a prerequisite artefact for multiple steps in the DSEEP standard for guiding design and development of distributed simulation environments (IEEE, 2022).

The Simulation Conceptual Modeling in the Design and Development of Simulations section explained the important role of the conceptual modeling process, and development of SCM documentation, in supporting the design and development of new simulations. The maintenance of SCM documentation is also critical for supporting the maintenance of simulations, enabling engineers to better understand implications of engineering changes. Simulation users rely on SCM documentation to inform their decisions regarding which simulation(s) to use to address a research question for experimentation or analysis, or to support training. The integration of multiple simulations into distributed simulation environments is a challenge that is largely unique to military simulation communities, and conceptual model documentation is necessary for informing decisions regarding semantic and conceptual interoperability of simulations.

The absence of SCM documentation among DoW simulations serves as a significant obstacle to compliance with the DoW directives for conducting VV&A activities in support of DoW simulations (DoD, 2024; Roca, 2013). Execution of VV&A, and verification in particular, depends on the availability of current conceptual model documentation. The U.S. Navy VV&A

Implementation Handbook provides thorough guidance for the execution of simulation V&V activities. A particularly notable distinction provided by this handbook is the distinction between two types of simulation verification. Requirements verification refers to the determination of how thoroughly the simulation aligns with the list of required capabilities specified in the requirements documentation. Design verification refers to the evaluation of how well the simulation aligns with the conceptual model documentation.

While both of these verification activities are important for the VV&A process, the design verification activity is particularly important for the M&S domain and depends on the availability of conceptual model documentation. Design verification helps ensure that all aspects of the intended simulation design are reflected in a simulation sufficiently to support its intended use. This is unique to M&S capabilities, and it may be easy for an acquisition organization to overlook the design verification requirement.

There are additional benefits of SCM documentation for supporting implementation of the Modular Open Systems Approach (MOSA) to simulation acquisition and management. MOSA is a well-documented acquisition and design approach that seeks to leverage existing standards in support of system use. MOSA couples technical and business architectures to facilitate components to be added or removed throughout a life cycle to help realize more efficient operation while encouraging innovative competition in system development (DoD, 2025a). The statute Title 10, U.S.C. 4401 requires MOSA for all DoD acquisition programs to the maximum extent possible. Recent guidance from the SECWAR highlights the importance of MOSA for supporting system interoperability and enabling "cost-effective and responsive modernization and sustainment of weapon systems" (DoD, 2025a). These benefits of MOSA apply equally to its implementation for M&S capabilities, and the directives correspond to a mandate for conceptual model development and maintenance for simulations. Complying with implementation of MOSA for M&S capabilities would require first ensuring the delivery and maintenance of the conceptual model documentation to inform consideration of system reuse and interoperability.

Shifting the Focus from Requirements to Simulation Development Process

The requirements transition process of the JCIDS acquisition process was not conducive to refinement of the simulation conceptual model documentation, which is critical in the early stages of simulation design and development. With the retiring of JCIDS and transition to the JFRP, the DoW has an opportunity to incorporate simulation design and development best practices in its acquisition process.

Modernizing the Simulation Requirements Transition Process

Simulation conceptual model documentation, which has heretofore been treated as a nice-to-have artifact in the simulation acquisition process, should be used as the linchpin in the transition of requirements specifications to the acquisition PM. Far from adding a bureaucratic obstacle, this focus on development of conceptual model documentation offers a tangible path for accelerating the acquisition process.

The conventional requirements generation and transition process breaks the iterative design and communication dynamics that are critical for the simulation design and development process. The simulation requirements and conceptual model transition process that we propose, depicted in Figure 5, fixes this by restoring iterative design and communications to the process by moving the final transition of the requirements (from the RM to the PM) until after an initial conceptual model has been developed and reviewed. This empowers the RM to ensure the simulation is appropriately designed to support the intended use. This process also provides the acquisition PM with a more value-added artifact, a blueprint for the simulation, to guide their consideration of available resources and engage with potential vendors for simulation development. Figure 5 details the proposed process flow, making explicit the iterative nature of the communications between the RM and the PM / developers in the comparison of the SCM against the requirements, and also clarifies that it is the RM who determines when the PM / developers have sufficiently understood the requirements and developed the SCM appropriately. Ultimately, the transition of the program to the acquisition PM takes place after the initial conceptual model documentation has been validated by the RM.

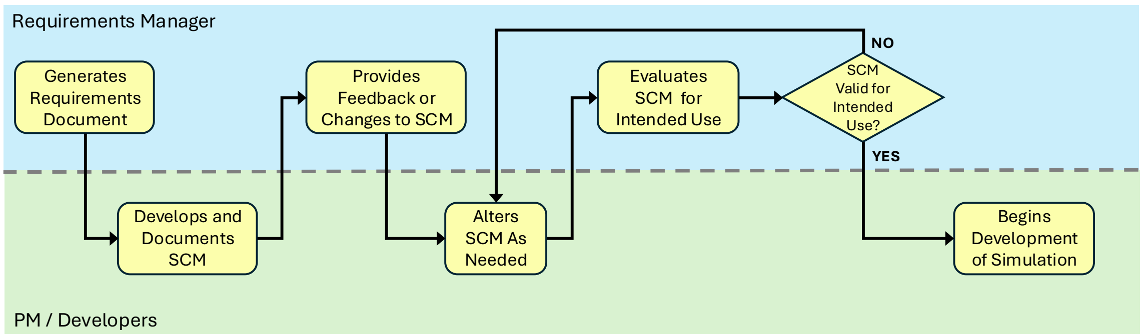

Figure 5. Proposed Simulation Requirements and Conceptual Model Transition Process

This proposed adjustment aligns with the intent behind the DoW transition from JCIDS to the JFRP, with the RM supporting the design of more tangible, value-added artifacts in the design of capability solutions. Rather than making such SCM development the sole responsibility of the PM, it would instead require increased collaboration between the RM and the acquisition PM's team, working together to refine conceptual model documentation to meet the needs of both parties.

M&S capabilities are unlikely to rise to the level of the Joint Operational Problems (JOPs) supported by the JFRP's Mission Engineering and Integration Activity (MEIA; JCS, 2026). That said, the JFRP provides services the autonomy to implement a similar approach, moving the requirements generation and transition process away from lengthy, bureaucratic requirements documents toward more value-added mission engineering artifacts. This presents an opportunity for increased collaboration between RMs and acquisition PMs in the iterative design and validation of conceptual model documentation for simulation capabilities. Successfully implementing this change to the requirements generation and transition process will require changes by both the RM and the acquisition PM.

The RM must go beyond writing a requirements document that lists required simulation capabilities. To facilitate the conceptual modeling process, the RM must also provide more thorough descriptions of the intended use(s) under which the simulation capability will be employed. In some cases, the RM will be describing existing environments where simulation capabilities are being replaced. More challenging situations, where explicit definition of intended use is even more important, are those where the RM will have to anticipate (or prescribe) how organizations' processes will change. The explicit definition of different intended use constructs will provide the baseline for the development of simulation conceptual model documentation (e.g., one set of SCM documents describes how the simulation will be used to support individual training; another set will illustrate its capabilities and limitations for an LVC staff training environment, etc.). Both the RM and the acquisition PM must share a common understanding of the simulation's intended use and participate in the refinement of the simulation conceptual model documents, to ensure these artifacts address the needs of both parties.

Proposed Process Alignment with VV&A and MOSA Requirements

In addition to facilitating a more effective requirements transition, the proposed simulation requirements and conceptual model transition process would provide guardrails for ensuring compliance with DoW VV&A regulations and achieving MOSA and reuse for M&S capabilities. The proposed process would ensure SCM documentation is available to support VV&A activities, even formalizing an early VV&A activity step that can easily be overlooked: conceptual model validation. By requiring the RM to validate that the SCM aligns with the intended use and requisite representation of the referent, the proposed process can ensure the conceptual model documentation is appropriately scoped to support simulation development and subsequent VV&A activities (Çilden et al., 2023; Tolk et al., 2013).

In supporting explicit SCM documentation and subsequent execution of VV&A activities, the formalization of conceptual model documentation delivery also supports another priority in DoW M&S management: reuse. More than just a cost saving measure, M&S capability reuse is critical due to the increasing demand for highly specialized simulations. Analysis, experimentation, and training communities rely on combinations of simulations with varying levels of detail and fidelity across different domains (e.g., electromagnetic warfare, space, and maritime operations). DoD Instruction 5000.61 identifies the need for VV&A standards to “foster reuse of DoD models, simulations, distributed simulations, and associated data.” Just as VV&A documentation is necessary for substantiating analysis findings, SCM documentation is necessary for providing different stakeholders with an understanding of (and confidence in) simulation capabilities to support M&S capability reuse.

By ensuring the delivery of SCM documentation, the proposed process also provides a path forward for enhanced compliance with DoW mandates for implementation of MOSA. That said, the narrow focus of this recommended process refinement would only address the initial conceptual model documentation generation and not the maintenance of that documentation that must be conducted throughout the simulation’s life cycle. To increase opportunities for reuse, conceptual model documentation should also be expanded as stakeholders find new intended uses. The reason for this is best illustrated through consideration of human-in-the-loop (HitL) simulations. Many HitL simulations are employed in different ways. The JCATS simulation can be used to train an individual servicemember who interacts directly with the simulation’s GUI or it can be used behind the scenes to support a large staff training exercise, stimulating command and control systems with direct interaction only occurring with response cell personnel. Developing a single set of conceptual model documentation for the simulation under these vastly different intended uses would be challenging if not impossible due to the need for the CM documentation to include the nature of user interactions (e.g., environmental cues provided to the user, affordances provided to the user for interactions with entities, feedback on user actions necessary for facilitating learning).

Overcoming Obstacles to Implementation

Despite the demonstrated value of conceptual modeling in the simulation development process, and requirements for conceptual model documentation throughout DoW simulations’ life cycles, DoW M&S practitioners face several obstacles in the implementation of these best practices. The term “conceptual model” is often misunderstood and misused, resulting in a perception among industry and government stakeholders that the practice itself is overly ambiguous, convoluted, and unnecessary. Simulation developers often perceive conceptual model documentation as a risk to their intellectual property and may be hesitant (or outright opposed) to providing it. These obstacles are largely due to a misunderstanding of simulation conceptual modeling aims and methodologies. While the DoW M&S community has long sought to overcome these obstacles (Pace, 2011), success is unlikely without clearly addressing the perceived obstacles and securing buy-in from the DoW acquisitions professionals.

Conceptual Model Documentation as Deliverables

Under the legacy JCIDS-based acquisitions process, SCM documentation was often viewed as a “nice-to-have,” low priority deliverable with limited impact on the requirements generation and transition process. The proposed process depicted in Figure 5 would change this situation, enabling RMs, acquisition professionals, and simulation developers to leverage simulation design and development best practices for enhanced communication and collaboration. Under this process, the SCM documentation is a critical tool that augments and enhances the design and development process to improve the alignment between simulation design and requirements, rather than an optional artefact to be generated after simulation development.

Simulation Conceptual Model Documentation Methodologies and Notations

A common refrain from industry and government acquisition professionals, when asked to provide simulation conceptual model documentation, is that there is insufficient guidance regarding conceptual modeling methodologies and standards for their implementation. Despite several papers, standards, and books providing guidance regarding simulation conceptual modeling approaches and best practices, there is no universal modeling notation or conceptual modeling methodology. The DoD Architecture Framework (DoDAF) provides a framework which may support the organization of some conceptual model documentation, but it does not provide guidance or standards for specific modeling notations to use for different simulation types or domains (van den Berg & Lutz, 2015).

In the 1990s, the DoD attempted to develop a comprehensive conceptual modeling methodology that could support documentation for many types of domains. This effort was abandoned, however, as the growing number and complexity of domains supported by modern simulations makes any such effort infeasible. Instead, guidance regarding the selection and application of conceptual modeling notations and methodologies must be developed for different communities, drawing on conceptual modeling best practices for developing SCM documentation that best supports critical simulation development and management activities (e.g., design and development, VV&A, database management).

Various researchers have provided frameworks to guide the development of simulation conceptual models and explain their role in the simulation design and development process. The conceptual modeling framework presented by Robinson (2008) provides guidance regarding various actions which should be taken throughout the conceptual modeling process, such as defining simulation inputs/outputs and documenting simulation simplifications and abstractions. Robinson briefly identifies some potentially useful diagrammatic modeling notations for simulation conceptual modeling (e.g., process flow diagrams, activity cycle diagrams, event graphs, simulation activity diagrams). Turnitsa and Tolk (2011) provide a taxonomy of conceptual models to provide additional context for consideration of different conceptual modeling techniques. The Multi Viewpoint Conceptual Modeling (MVCM) methodology presents lower-level guidance for how different notations and viewpoints may be combined in a conceptual modeling methodology approach that is tailorable to support broad use across different domains (Morse & Drake, 2020).

When considering which conceptual modeling methodologies or notations to employ, it is critical for all stakeholders to remember the purpose of the documentation. In the acquisition process, the SCM documentation is intended to support communication between the stakeholders. This communication will be for different purposes, at times requiring greater or lesser detail for different types of content (Aros, 2025). Different kinds of information may be best communicated using different conceptual modeling documentation methodologies. No one specific modeling notation will be sufficient. Instead, stakeholders must collaborate to identify the SCM documentation methodologies that best capture simulation design considerations and make them accessible to appropriate stakeholders. Identification of appropriate modeling notations and methodologies for reuse in support of different domains will also yield benefits in the comparison of simulations for reuse and interoperability considerations.

Developer Reluctance to Share Simulation Design Documentation

A primary hesitancy expressed by some simulation developers and acquisition professionals relates to the disclosure of proprietary information. This may reflect a misunderstanding of conceptual modeling; simulation conceptual models are implementation agnostic, so they don't provide insights into the details regarding how the simulation code is structured or implemented. That said, in some situations there may be a legitimate concern regarding proprietary design elements detailed in the SCM, such as unique logic or algorithms that necessitate protection from competitors. In these situations, conceptual model documentation can be protected through appropriate controlled unclassified information designation and handling procedures.

Some vendors are also reluctant to provide insight into simulations' limitations. The illumination of simulation limitations should not be considered to be a detriment. All simulations consist of abstractions, and are inherently limited in their representation of the referent in some way. It may appear counterintuitive to highlight the limitations of one's simulation's ability to represent a domain; however, providing this insight will likely make the simulation more useful to analysts and training designers who may consider leveraging it. If the limitations of a simulation are unknown to an analyst who is designing a study, they incur an unknown amount of risk if they employ the simulation. On the other hand, if an analyst has a less powerful simulation available, but they understand the nature of the simulation's abstractions, the analyst may be able to design a study to account for the simulation's abstractions, providing higher confidence in the results.

Discussion and Conclusions

The U.S. military M&S community has long recognized a need for reform, to increase compliance with industry best practices for simulation design and development. Limited employment of conceptual modeling practices presents a risk to simulation development programs. Limited availability of conceptual model documentation for existing M&S capabilities degrades the U.S. military's ability to effectively employ, manage, and reuse these capabilities. Despite multiple efforts by DoW M&S leadership and practitioners to increase the employment of conceptual modeling, the fundamental structure of the legacy requirements generation and transition process disincentivized use of conceptual model documentation. Reform is needed to ensure acquisition professionals are provided clear guidance for how to meet the intended use. Requirements managers must also be empowered to collaborate on the development of conceptual models rather than having to extensively document all requirements in an effort to provide complete clarity before the requirements hand-off. With the end of JCIDS comes an opportunity to implement industry best practices in the acquisition process for M&S capabilities. The proposed simulation requirements and conceptual model transition process presents a way to capitalize on this opportunity and to restructure the simulation acquisition process to leverage industry best practices by formalizing the role of simulation conceptual model documentation.

References

Abdelmegid, M. A., O'Sullivan, M., Gonzalez, V. A., Walker, C. G., & Poshdar, M. (2022). A case study on the use of a conceptual modeling framework for construction simulation.

SIMULATION: Transactions of the Society for Modeling and Simulation International, *98*(5), 433–460.

Appleget, J., Blais, C., Burks, R., Brown, R. F., Duong, D., Jaye, M., Perkins, T., & Thompson, M. (2011). *Irregular warfare model validation best practices guide*. TRADOC Analysis Center.

Aros, S. K. (2025). Content considerations for simulation conceptual model documentation. In *Proceedings of the 2025 Winter Simulation Conference (WSC)* (pp. 2193–2203). IEEE.

Borah, J. (2003). *SISO conceptual modeling tutorial (Full)*. Presented to the SISO Simulation Conceptual Modeling Study Group.

Brooks R. J. & Wang W. (2015). Conceptual modelling and the project process in real simulation projects: A survey of simulation modellers. *Journal of the Operational Research Society, 66*(10), 1669–1685.

Çilden, E., Sezer, A., Canberi, M. H., & Oğuztüzün, H. (2023). *Iterative and incremental validation of simulation conceptual models.* 2023 Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC), Orlando FL, United States.

DoD. (2020). *Operation of the software acquisition pathway* (DoDI 5000.87).

DoD. (2024). *DoD modeling and simulation verification, validation, and accreditation* (DoDI 5000.61).

DoD. (2025a, February). *Implementing a modular open systems approach in Department of Defense programs*.

DoD. (2025b, November). *Transforming the defense acquisition system into the warfighting acquisition system to accelerate fielding of urgently needed capabilities to our warriors*.

IEEE Computer Society. (2022). *IEEE recommended practice for distributed simulation engineering and execution process (DSEEP)* (IEEE Std 1730–2022).

Joint Chiefs of Staff. (2021, October 30). *Manual for the operation of the Joint Capabilities Integration and Development System* (JCIDS). https://www.waru.edu/sites/default/files/2024-01/Manual%20-%20JCIDS%20Oct%202021.pdf

Joint Chiefs of Staff. (2026, January 15). *Manual for the Joint Requirements Oversight Council and the Joint Force requirements process* (CJCSM 5123.01). https://www.jcs.mil/Portals/36/Documents/Library/Manuals/CJCSM%205123.01.pdf

Morse, K. L. & Drake, D. L. (2020, September 22). *Multi-viewpoint conceptual modeling: In support of simulation interoperability readiness levels (SIRLs)* [Conference Presentation]. NATO CA2X2 Forum 2020, Virtual Event.

NATO (2012). *Conceptual modeling (CM) for military modeling and simulation (M&S): Final report of MSG-058* (RTO TECHNICAL REPORT TR-MSG-058). https://apps.dtic.mil/sti/pdfs/ADA569241.pdf

Pace, D. K. (2000, October). Conceptual model development for C4ISR simulations. In *Proceedings of the 5th international command and control research and technology symposium* (pp. 24–26). CCRP Press.

Pace, D. K. (2011). Conceptual modeling evolution within US Defense communities: The view from the simulation interoperability workshop. In S. Robinson, R. J. Brooks, K. Kotiadis and J. van der Zee (Eds.), *Conceptual modeling for discrete-event simulation* (pp. 423–449). CRC Press.

Petty, M. D. (2009). Verification and Validation. In J. A. Sokolowski & C. M. Banks (Eds.), *Principles of modeling and simulation: A multidisciplinary approach* (pp. 121–148). Wiley.

Robinson, S. (2008). Conceptual modeling for simulation part II: A framework for conceptual modeling. *Journal of the Operational Research Society*, *59*(3), 291–304.

Robinson, S. (2012). Tutorial: Choosing what to model—Conceptual modeling for simulation. In *Proceedings of the 2012 Winter Simulation Conference*.

Robinson, S., Arbez, G., Birta, L. G., Tolk, A., & Wagner, G. (2015). Conceptual modeling: definition, purpose and benefits. In *Proceedings of the 2015 Winter Simulation Conference* (pp. 2812–2826). IEEE.

Roca, R. (2013). M&S conceptual modeling as an enabler of M&S reuse, agile, and open-source: A TRS initiative in response to the M&S SC priority objectives. *M&S Journal, 8(1),* 20–27.

Tolk, A., Diallo, S. Y., Padilla, J. J., & Herencia-Zapana, H. (2013). Reference modelling in support of M&S—Foundations and applications. *Journal of Simulation*, *7*(2), 69–82.

Turnitsa, C. D., Tolk, A. (2011). *A taxonomy of conceptual models* (Paper 11S-SIW-039). 2011 Simulation Interoperability Workshop. Simulation International Standards Organization.

van den Berg, T. W., & Luz, R. (2015). *Simulation environment architecture development using the DoDAF* (Paper 15F-SIW-019). 2015 Fall Simulation Interoperability Workshop. Simulation International Standards Organization.

Operationalizing Cyber Survivability Through Requirements Decomposition: A Marine Corps Case Study

Kathleen B. Coen—is a Cyber Test Engineer at the Marine Corps Tactical Systems Support Activity, where she leads adversarial cyber assessments of USMC weapon systems to identify mission-impacting vulnerabilities prior to fielding. She holds a BS in Computer Science from Georgetown University and an ScM in Cybersecurity from Brown University. She is currently a PhD student in Information Sciences at the Naval Postgraduate School. Her work focuses on cyber survivability, test and evaluation, and the integration of cybersecurity policy into acquisition and operational decision-making. She holds the CISSP and DAWIA Engineering & Technical Management Practitioner certifications. [kathleen.coen@usmc.mil]

Capt Andrew S. Benn, USMC—is an active-duty Marine Corps officer. A 2023 graduate of the Naval Postgraduate School's Computer Science program and a 2017 graduate of the U.S. Naval Academy Cyber Operations Program, he is currently serving as the Cyber Branch Head at Marine Corps Tactical Systems Support Activity. In this role he supports Marine Corps acquisitions programs with Cyber Developmental Test and Evaluation and integration with Defensive Cyber tools. He served previously with Combat Logistics Regiment 35 and 3d Maintenance Battalion in Okinawa, Japan, as a communications officer and company commander. [andrew.benn@usmc.mil]

Abstract

United States Marine Corps (USMC) warfighting systems must operate reliably in contested cyber environments to bring their capabilities to bear to a future fight. To ensure mission-critical and safety-critical functions and components remain operational, these systems must be equipped to prevent, mitigate, recover from, and adapt to adverse cyber events: a concept known as cyber survivability. Joint Staff J6 defines 10 Cyber Survivability Attributes (CSAs) and provides guidance for requirements and resource sponsors early in the acquisition lifecycle. However, as programs progress through the acquisition lifecycle, there exists no additional guidance on how to derive, validate, and verify standardized and measurable system-level cyber survivability requirements. This lack of traceability between policy, security controls, and system engineering artifacts results in inconsistent implementation, redundant testing, and reduced ability to evaluate survivability. The Marine Corps Tactical Systems Support Activity's (MCTSSA's) Cyber Branch worked with various USMC Program Offices and other stakeholders to decompose cyber survivability requirements into tailored performance specifications, verification processes, and acceptance criteria. Through iterative application to multiple USMC Programs of Record, this design science research led to the service's inaugural guidance on cyber survivability within the Warfighting Acquisition System: the USMC Cyber Survivability Requirements Guidebook.

Keywords: cyber survivability, cybersecurity, requirements decomposition, test and evaluation

Introduction and Background

Warfighting systems must operate reliably in contested cyber environments to bring their capabilities to bear to a future fight, whether the conflict is with an asymmetric adversary or a pacing threat. To ensure mission-critical and safety-critical functions and components remain operational, systems must be equipped to prevent, mitigate, recover from, and adapt to adverse cyber events (DoW, 2026). This is a concept known as *cyber survivability*. Cyber survivability is critical to the success of modern warfighter systems, especially as adversaries increasingly target these systems through cyberattacks.

Since its introduction in 2015, cyber survivability has been included as a system requirement via the System Survivability Key Performance Parameter (SS KPP; Pitcher, 2023). The SS KPP is intended to promote the development of critical warfighter capabilities that can survive kinetic attacks and non-kinetic threats across domains and applicable environments. cyber survivability is one element of the SS KPP.

Striving for cyber survivability requires a balanced approach that integrates cybersecurity and cyber resilience. Cybersecurity is the process of protecting systems from unauthorized access and malicious attacks by ensuring the confidentiality, integrity, and availability of data and services. In contrast, cyber resilience assumes that security measures may eventually fail and emphasizes the system's ability to operate under duress, recover, and maintain operational capabilities (Swanson et al., 2010). Cybersecurity and cyber resilience form cyber survivability, ensuring systems can withstand and recover from a wide range of threats.

Cyber survivability is accomplished through a combination of Cyber Survivability Attributes (CSAs). The CSAs (Table 1) are a holistic set of cybersecurity and cyber resilience requirements. Each CSA corresponds to one of the four SS KPP pillars (prevent, mitigate, recover, adapt). By incorporating the concepts of cyber resilience, CSAs not only support but go beyond the coverage provided by the Risk Management Framework.

Table 10. Cyber Survivability Attributes

SS KPP Pillars	Cyber Survivability Attributes (CSAs)
Prevent	CSA-01: Control Access
	CSA-02: Reduce Cyber Detectability
	CSA-03: Protect Data in Transit
	CSA-04: Protect Data at Rest
	CSA-05: Protect Critical Functions
	CSA-06: Minimize and Harden Attack Surfaces
Mitigate	CSA-07: Baseline and Monitor Systems to Detect Anomalies
	CSA-08: Enable Cyber Defense
Recover	CSA-09: Recover Capabilities
Adapt	CSA-10: Sustain an Operationally Relevant Cyber Survivability Risk Posture (CSRP)

The CSAs are primarily intended for requirements and resource sponsors developing an Initial Capabilities Document (ICD), Capability Development Document (CDD), Information System ICD (IS-ICD), or Information System CDD (IS-CDD; Joint Staff, 2025).

As programs progress through the acquisition lifecycle, they transition from their service-level requirements process (replacing the Joint Capabilities Integration and Development System [JCIDS] process) to the Warfighting Acquisitions System (WAS). Within WAS, high-level requirements should be decomposed into performance specifications for engineering and testing purposes.

However, at the start of this research effort, there existed no additional guidance on how to derive, validate, and verify standardized and measurable system- and component-level cyber survivability requirements. This lack of traceability between policy, security controls, and system engineering artifacts resulted in inconsistent implementation, redundant testing, and reduced ability to evaluate survivability.

This gap has direct operational and acquisition consequences. As the Government Accountability Office (GAO) found, incorporating cybersecurity practices from the earliest stages of acquisition is typically easier, less costly, and more effective than attempting to add, or "bolt on," protections later in the development cycle or after a system is fielded. When cyber survivability requirements are not clearly defined and translated into measurable system-level performance specifications, acquisition programs are limited in their ability to assess whether systems can continue to perform mission-critical functions under contested cyber conditions. This deficiency undermines risk-informed acquisition decisions, weakens the effectiveness of developmental and operational testing, and increases the likelihood that systems are fielded with unresolved vulnerabilities that could disrupt mission execution (GAO, 2021).

This research effort addresses this problem by providing analysis and guidance on cyber survivability requirements decomposition, leveraging a systems engineering process to translate high-level requirements into measurable and verifiable specifications. The central research question is:

> How can a structured requirements decomposition framework be used to operationalize CSAs into measurable and testable acquisition artifacts that improve test coverage and support risk-informed decision-making across the acquisition lifecycle?

Methods

This study leveraged a design science research (DSR) approach. DSR is an established approach in information systems and engineering disciplines for developing and evaluating artifacts intended to solve identified organizational problems (Hevner et al., 2004). The primary contribution lies in both the artifact itself and the evidence demonstrating its utility. A design science approach was chosen because it places emphasis on clarifying the problem statement, goals, and underlying theoretical constructs to generate a new artifact (Gregor & Jones, 2007; McLaren et al., 2011; Moon & Ngai, 2010). DSR enables the development of structured, repeatable solutions that can be applied and evaluated within operational environments.

The DSR process follows an iterative cycle consisting of (1) problem identification, (2) definition of solution objectives, (3) artifact design and development, (4) demonstration in a relevant context, and (5) evaluation of artifact effectiveness (Peffers et al., 2007). This study adheres to that structure, using the iterative development and ultimate application of a requirements decomposition framework to operationalize cyber survivability.

Problem Identification

The Marine Corps Tactical Systems Support Activity (MCTSSA) Cyber Branch provides cyber test and evaluation (T&E) support to Marine Corps Program Offices. During these engagements, system documentation, such as Capability Development Documents (CDDs), System Requirements Documents (SRDs), and Requirements Traceability Matrices (RTMs), is reviewed and analyzed to inform test activities.

Across multiple programs, a consistent pattern was observed. While CSAs were identified within the CDDs, there was limited evidence that they were decomposed into system- and component-level performance specifications. GAO (2021) found numerous examples of government-generated documentation omitting cyber-related performance specifications cybersecurity requirements. If included, they were often not measurable or directly verifiable through test activities.

This disconnect indicated a lack of standardized mechanisms for translating survivability constructs into actionable engineering and test artifacts, resulting in inconsistent implementation and reduced evaluability of cyber survivability.

Definition of Solution Objectives

From September 2024 to April 2025, MCTSSA Cyber held a series of meetings with various stakeholders from several Program Offices from Marine Corps Systems Command and Program Executive Office–Land Systems, Systems Engineering and Acquisition Logistics (SE&AL), other USMC T&E Activities, and representatives from the Office of the Under

Secretary of War for Research and Engineering (OUSW[R&E]). Meetings with Program Offices included participation from Information Systems Security Managers (ISSMs), Information Systems Security Engineers (ISSEs), and Information Systems Security Officers (ISSOs) that are responsible for the implementation of CSAs. Meetings were either held virtually (via Microsoft Teams), in person at Marine Corps Base Camp Pendleton, CA, or in person at Marine Corps Base Quantico, VA, and typically lasted 1 hour.

Initial meetings were used to identify and explore the problem and understand stakeholders' perspectives. The MCTSSA Cyber Branch incorporated this feedback into initial drafts of the artifacts, which were then presented, discussed, and refined, in an iterative fashion, at future meetings.

The iterative development and application of the artifact yielded several emergent design principles for operationalizing cyber survivability within acquisition environments. First, survivability constructs must be decomposed into traceable, tailorable performance specifications that preserve alignment with policy-level intent while allowing system-specific refinement. Second, verification methods and plans must remain Tool, Technique, and Procedure (TTP)–agnostic but requirements-specific, enabling consistent test coverage without constraining T&E organizations' testing creativity. Third, artifact adoption depends on integration with existing acquisition workflows and taskings, rather than introducing new processes or burdens. These findings extend prior design science work by proposing a reusable set of principles for translating abstract cybersecurity constructs into actionable acquisition artifacts.

As a result, the MCTSSA Cyber Branch defined a set of objectives for a structured solution to support survivability operationalization within acquisition programs. The solution was required to

- facilitate the decomposition of CSAs into measurable, system-level performance specifications;
- align derived requirements with National Institute of Standards and Technology Special Publication (NIST SP) 800-53 controls, the backbone of the DoW's Risk Management Framework;
- provide associated verification methods and plans to support developmental and operational cyber T&E;
- support said evaluation of implementation effectiveness by providing defined metrics; and
- integrate with existing DoW acquisition and engineering processes without introducing additional procedural burden.

The solution was intended to serve both Program Offices, responsible for requirements development and system acquisition, and T&E activities, responsible for verification and validation of cyber survivability.

Artifact Design and Development: The Marine Corps Cyber Survivability Requirements Guidebook

The ultimate artifact from this research is the *United States Marine Corps Cyber Survivability Requirements Guidebook*. This guidebook offers a structured approach for defining and evaluating the 10 CSAs critical for maintaining Cyber survivability within USMC systems. Each CSA section is accompanied by a summary of the CSA and its specific objectives, followed by a comprehensive listing of requirements in support of the CSA. Each requirement includes exemplar language for the performance specification, measurable metrics, and verification methods to support system evaluations during both developmental and operational cyber testing phases. This format aims to provide Marine Corps acquisition activities and the

Test and Evaluation community with a clear, actionable framework for CSA refinement and evaluation.

1. CSA Summary

Each CSA section begins with a summary of the CSA describing the attribute's purpose and strategic significance. The summary highlights the primary objective of the CSA, key enablers that contribute to achieving it, and the anticipated mission impact of effective implementation. This provides a holistic understanding of how each CSA supports system resilience and mission continuity.

2. Performance Specifications

Each CSA section includes a listing of performance specifications in direct support of that CSA. Performance specifications provide subsystem- and component-level requirement statements that guide system implementation. As standard with performance specifications, functional requirements state the required results without specifically stating how the results are to be achieved; the guidebook does not present a preconceived solution to each requirement.

Performance specifications listed in the guidebook should be further tailored for each system. Provided statements do not include behavior considerations, such as constraints and restraints, to leave the trade-space open for program managers. An example of tailoring is shown in ***Table 11***.

Table 11. Example of Performance Specification Tailoring

Performance Specification Consideration	Guidebook's Tailorable Performance Specification	System-Specific Tailored Performance Specification
Cross Domain Solutions (CDSs)	The system shall integrate Department of Defense (DoD)-approved CDSs to enable data access between different security domains, per DoD Instruction (DoDI) 8540.01, May 8, 2015, Incorporating Change 1, August 28, 2017.	The system shall integrate DoD-approved CDSs, to including data labeling and auditing, per DoDI 8540.01, May 8, 2015, Incorporating Change 1, August 28, 2017, to enable data access from Coalition to Secret domains in the data standards or file formats of Extensible Markup Language (XML) and Internet Relay Chat (IRC).
Automated Shedding on Non-Mission and Non-Safety Critical Functions	The system shall provide capabilities to shed non-mission and non-safety critical functions, systems/subsystems, and interfaces.	The system shall provide automated capabilities to shed non-mission and non-safety critical functions, systems/subsystems, and interfaces within 5 seconds of detection of a critical system overload or a verified cyberattack. Non-mission and non-safety critical functions are defined in Table X. Shedding shall occur in the order specified in Table Y, minimizing impact to overall system performance as detailed in Table Z. Shed functions shall be automatically re-enabled when system load returns to normal operating parameters.

Performance specifications should be selected during the Technology Maturation and Risk Reduction phase and finalized at the Critical Design Review that takes place during the Engineering and Manufacturing Development phase. Performance specifications should be included in the System Requirements Document or the System Performance Specification (SPS). The SPS is included in the Request for Proposal package provided to contractors and is closely aligned with the program's CDD and Capability Production Document, ensuring ultimate alignment with the cyber element of the SS KPP.

3. Verification Methods and Plans

For each performance specification, the guidebook provides a listing of verification methods and plans to test implementation effectiveness. The verification plans provide guidance on what should be tested but does not specify how they should be tested; verification plans were purposely left non-prescriptive. The T&E community should retain full autonomy over specific tools, techniques, and procedures used during testing. The ultimate objective is consistent and thorough assessments, ensuring test plans and cases have adequate coverage for each requirement.

Each verification plan was tagged with its associated verification method. As stated in the *Systems Engineering Guidebook*, verification plans can be achieved through any combination of the following methods:

- **Inspection or Examination:** Visual inspection of equipment and evaluation of drawings and other pertinent design data and processes should be used to verify conformance with characteristics, such as physical, material, part, and product marking and workmanship.
- **Demonstration:** Demonstration is the performance of operations at the system or system element level where visual observations are the primary means of verification. Demonstration is used when quantitative assurance is not required for the verification of the requirements.
- **Analysis:** Analysis is the use of recognized analytic techniques (including computer models) to interpret or explain the behavior/performance of the system element. Analysis of test data or review and analysis of design data should be used as appropriate to verify requirements.
- **Test:** Test is an activity designed to provide data on functional features and equipment operation under fully controlled and traceable conditions. The data is subsequently used to evaluate quantitative characteristics.

In accordance with *the Guide for Performance Specifications*, each performance specification can, and typically does, utilize more than one method for proper verification.

The verification methods and plans should be used as a starting point for the development of test plans and cases. Cyber test activities must begin as soon as the program is established and continue throughout the acquisition lifecycle. USMC cyber T&E activities and teams should be involved as soon as possible in the early system assessments based on available program documentation, hands-on testing of individual components/subcomponents during system development, and cyber survivability testing of complete systems and the entire platform in test and operationally representative environments.

4. Metrics

Each performance specification is accompanied by a listing of related metrics used to determine the successful implementation of the requirement. These metrics can be used to determine if the system's specific threshold survivability requirements are being met and to quickly gauge a system's potential performance under contested cyber conditions. These acceptable threshold requirements are not numeric in the guidebook, as they are highly system- and mission-dependent and should be determined by the PM.

Through this structure, the guidebook provides a repeatable mechanism for transforming abstract survivability constructs into actionable engineering and acquisition artifacts. In doing so, it addresses the core problem identified in this research: the lack of standardized methods for operationalizing cyber survivability within the acquisition lifecycle.

Results: Demonstration in a Relevant Context

Consistent with design science methodology, the artifact was not evaluated in isolation but via application within an operational context.

A draft version of the guidebook was used by a Marine Corps Program of Record within a Middle Tier Acquisition pathway, where it was used to develop and refine cyber survivability requirements and associated verification methods.

During the requirements phase, the platform was assigned a Cyber Survivability Risk Category (CSRC) value of 3. The CSRC identifies the appropriate degree of cyber survivability required for a system. It is a function of the system's Mission Type, the Adversary Threat Tier expected to be facing the system, the Cyber Dependence Level of the system, and the Impact Level of system compromise or loss; the final CSRC is a numeric value ranging from 0 to 5. The CSRC value determines how many CSAs should be selected; a CSRC 3 system does not require all 10 CSAs. This platform selected seven CSAs total.

Evaluation of Artifact Effectiveness

In the end, based on the seven CSAs, the platform was assigned a robust set of cyber-related performance specifications. 90% of the performance specifications were taken and tailored from the guidebook. The program averaged eight verification methods per performance specification. 82% of the verification methods were taken and tailored from the guidebook. The guidebook is meant to provide initial considerations that should be built on and further tailored for each system.

Feedback from Program Offices indicated increased confidence in cyber requirements alignment, reduced ambiguity in CSA interpretation, and improved communication between the dedicated professionals involved in requirements development, management, implementation, and evaluation of cyber survivability in the USMC.

The artifact of this DSR, the *USMC Cyber Survivability Requirements Guidebook*, is expected to be published June 2026.

Discussion

The results of this work provide empirical support for the central research question that guidance for a structured requirements decomposition process improves the operationalization of cyber survivability within acquisition programs. By translating high-level CSAs into measurable and testable performance specifications, the artifact addresses and minimizes the critical gap between strategy intent and engineering execution.

Cyber survivability, as it is defined in high-level guidance, lacks measurability. Without further decomposition, the CSAs remain a conceptual objective rather than an engineering requirement. The artifact effectively transforms CSAs into a set of verifiable system design consideration, enabling engineers to design against concrete criteria and testers to evaluate performance in a structured manner.

The study also highlights the role of decomposition in improving traceability across the acquisition lifecycle. By establishing explicit linkages between CSAs, NIST SP 800-53 controls, Zero Trust Activities, and system-level requirements, the framework created a continuous thread for those in the Systems Security Engineering community to focus on. In the absence of such linkages, acquisition programs risk redundant engineering and cyber assessments.

From an acquisition standpoint, the implications are equally significant. The introduction of structured, decomposed requirements improved alignment across stakeholders, including Program Offices, system engineers, vendors, and test organizations. Reduced ambiguity in requirement interpretation will facilitate more consistent implementation, ease of understanding, and efficient communication. This is particularly important in complex acquisition environments, where misalignment between stakeholders, particularly between the government and its vendors, can lead to costly rework, delayed testing, and incomplete risk assessments.

However, the results also underscore that decomposition is not a purely technical solution. Its effectiveness depends on adoption, correct application, and integration within existing acquisition processes. The guidebook provides a structured starting point, but it requires tailoring to system-specific contexts and active engagement from stakeholders. As such, the value of decomposition lies not only in the artifact itself, but in how it is leveraged and applied by the broader engineering and acquisition community.

By bridging the gap between policy and implementation, the approach improves traceability, enhances test coverage, and supports more informed acquisition decisions. These findings contribute to ongoing efforts to integrate cybersecurity and resilience considerations into systems engineering and acquisition practices and provide a foundation for scaling survivability-focused approaches across the DoW.

Future Research

This study is subject to several limitations that inform both the interpretation of findings and the direction of future research. First, while several Program Offices and stakeholders provided feedback in the iterative cycle of artifact development, the evaluation of the artifact was limited to a single-case application within a Marine Corps platform. While the selected case provides a realistic and operationally relevant test environment, it limits the generalizability of findings across different acquisition pathways, system types, services, and organizational contexts. Programs with varying levels of cyber maturity, system complexity, or stakeholder expertise may experience different outcomes when applying a requirements decomposition framework. Additionally, the platform selected is still early in the acquisitions lifecycle. The long-term impact of requirements decomposition on ultimate system survivability and lifecycle cost has not yet been observed.

Finally, this research opens an opportunity to extend the decomposition approach beyond cyber survivability to other elements of the SS KPP or other high-level requirements. Each of these domains faces similar challenges in translating high-level constructs into actionable engineering requirements. A unified decomposition methodology may therefore provide broader value across the WAS.

References

DoW. (2026). *Cyber developmental test and evaluation* (DoW Manual 5000.103). https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodm/5000103m.pdf

GAO. (2021). *Weapon systems cybersecurity: Guidance would help DOD programs better communicate requirements to contractors* (GAO-21-179). https://www.gao.gov/assets/gao-21-179.pdf

Gregor, S., & Jones, D. (2007). The anatomy of a design theory. *Journal of the Association for Information Systems*, 8(5), 312–335. https://doi.org/10.17705/1jais.00129

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, *28*(1). https://doi.org/10.2307/25148625

Joint Staff J6 Command, Control, Communications, and Computers/Cyber. (2025). *Cyber survivability endorsement implementation guide*.

McLaren, T. S., Head, M., Yuan, Y., & Chan, Y. (2011). A multi-level model for measuring fit between a firm's competitive strategies and information systems capabilities. *MIS Quarterly*, *35*(4). https://doi.org/10.2307/41409966

Moon, K. L., & Ngai, E. (2010). R&D framework for an intelligent fabric sample management system: A design science approach. *International Journal of Operations Production Management*, *30*(7). https://doi.org/10.1108/01443571011057317

Peffers, K., Tuunanen, T., Rothenberger, M., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3). https://doi.org/10.2753/MIS0742-1222240302

Pitcher, S. (2023). Improving cyber survivability for weapon system mission assurance. *Cybersecurity and Information Systems Information Analysis Center*. https://csiac.dtic.mil/webinars/improving-cyber-survivability-for-weapon-systems-mission-assurance/

Swanson, M., Bowen, P., Phillips, A., Gallup, D., & Lynes, D. (2010). *Contingency planning guide for federal information systems* (NIST Special Publication No. 800-34). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-34r1

Acquisition Software Management Planning (SMP) Is Critical to Mission Success

Carol Woody—is a principal researcher in the CERT Division of the Software Engineering Institute (SEI) at Carnegie Mellon University. She leads a team that builds capabilities and competencies for measuring, managing, and sustaining software assurance and cybersecurity for highly complex software-intensive systems and supply chains throughout the Acquisition life cycle. She coauthored the book Cyber Security Engineering: A Practical Approach for Systems and Software Assurance, which was published by Pearson Education as part of the SEI Series in Software Engineering. The CERT Cybersecurity Engineering and Software Assurance Professional Certificate is based on the research she led. [cwoody@cert.org]

Mike Bandor—is a senior software engineer in the CERT Division of the Software Engineering Institute (SEI) at Carnegie Mellon University. He leads teams that enable organizations to enhance the predictable performance and mission assurance in the acquisition, evolution, and operations of software-reliant systems. He has more than 36 years of experience with military systems, including business systems, command and control systems, satellite systems, ground systems, aircraft, and ground-based radar systems. Prior to joining the SEI in 2005, he was an enlisted computer programmer in the U.S. Air Force, and he retired with almost 23 years on active duty. [mbandor@sei.cmu.edu]

Abstract

Today's systems are increasingly software intensive, complex, and reliant on third-party technology. We live in a world of systems of systems linked by software that connects services and hardware and essentially removes many previous human and geographic restrictions. Unfortunately, acquisition practices have not kept pace with these changes. Leadership is still primarily monitoring cost and schedule. Today's systems can be assembled faster and cheaper because software is rarely built for its intended use. Instead, much of it is reused, sourced from third parties (and increasingly from open source sites), but with increased risk. All software contains potential vulnerabilities that increase the risk of experiencing successful cyber attacks. It is critical to ensure that system requirements are met without extraneous behaviors that would jeopardize the mission. This paper explains why effective software management is critical to the acquisition of today's systems, which are primarily software intensive. It also shares lessons learned in current efforts underway to build and implement a Software Management Plan (SMP) in major Department of War (DoW) acquisitions and describes the research underway to improve how software is monitored and managed.

Overview of the Software Challenges

The use of software in acquisition is growing exponentially both as a component within every technology acquisition and as the enabler for creating software components. Software is also used as an enabler to connect legacy systems with new capabilities to enhance existing operational processes. Software is fast becoming the primary enabler of the functionality that systems provide. However, for most acquisition programs, software is "behind the scenes" facilitating a wide range of activities that were previously handled by hardware but with greater flexibility and lower cost.

Some of this software is built as part of the acquisition, but much of it is reused from similar platforms, brought into each software component from existing language libraries, or pulled from open source sites. For the prime contractor, only a finite portion of the software in the acquisition is built for purpose. Much of it is repurposed from other similar projects, acquired from subcontractors, linked into a system from specialized service providers (e.g., Cloud), or downloaded from open source sites. The acquirer must then confirm that these pieces that come from a range of sources and built by a range of development processes meet the intent of the contract governing the acquisition. An example of software component composition is shown in ***Figure 30***.

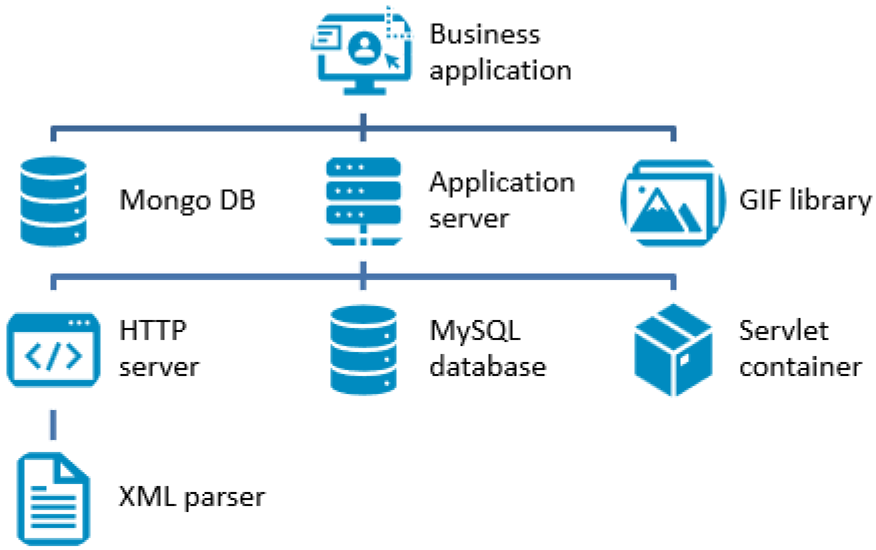

Figure 30.Assembly of Software from Components

To deliver systems that operate successfully in today's contested environments, which is becoming the norm for the Department of War (DoW), software must be assembled and maintained with high assurance. This result does not come automatically. Instead, it requires extensive planning and monitoring. All software contains potential vulnerabilities that increase the risk of successful cyber attacks (Woody et al., 2014). Therefore, it is critical to ensure that system requirements are met without extraneous behaviors that would allow attackers to jeopardize the mission.

There is a wide range of recent guidance that focuses on improving the speed of software delivery, such as DevSecOps guidance (DoD, 2024) and Software Acquisition Pathway guidance (DoD, 2020). In addition, there is recent guidance for program protection that identifies and addresses risks that software and the software supply chain represent to the mission (DoW, 2025). However, connecting this guidance to each acquisition is a challenge. While there is a great deal of software within every acquisition, it is spread somewhat invisibly throughout the system.

The requirements that are prepared for each acquisition are allocated to various system components, some of which may be outsourced. Each component is decomposed into subcomponents, which are composed of hardware, software, and firmware elements, that will be built by different processes and integrated to deliver the planned functionality. This decomposition is reflected in the Work Breakdown Structure (WBS) for the acquisition. If the component is outsourced, integration occurs first at the subcomponent level and then at the component level. For major systems, the link to software does not appear until levels 4 and 5 in the decomposition and WBS, and this link is not always clearly defined until Milestone B (DoD, 2022).

The primary focus of acquisition program management is typically on cost and schedule for the delivery of a well-engineered system that meets the mission. However, with software, it is possible to have a system that functions as intended but is riddled with vulnerabilities, which make it too easy to attack (Woody et al., 2026). The challenge in managing this software is twofold:

- Potential vulnerabilities can be introduced at each step in the development life cycle, and delays in identifying and removing these vulnerabilities will result in higher costs and potential schedule delays as the software is integrated into larger pieces of the system.
- Managing these many pieces of software is impossible without consistent standards that apply across the system. The acquisition process needs to include good software engineering processes and practices that will reduce the introduction of vulnerabilities and steps for managing software to ensure it conforms to the acceptable level of quality.

There are many ways to address these two challenges within an acquisition. The keys are to have a plan for how they will be managed and the resources with the appropriate knowledge to execute the plan across the life cycle. Moving into software development without a plan for how it will be managed from an acquisition perspective and confirmation that the vendors are capable of building and delivering software at the acceptable level of rigor puts the program at high risk. Too many acquisitions focus only on obtaining Software Bill of Material (SBOMs) to meet a compliance mandate (Wallen, 2023), but these will not provide a way of determining where potential vulnerabilities are located and their potential risk to the program. A Software Management Plan (SMP) is needed that is built on solid software management and assurance practices that are well integrated into the system acquisition to ensure that the software is acquired, integrated, tested, and implemented successfully to meet mission requirements.

An SMP is not the same as a Software Development Plan (SDP). An SDP is typically written by the developing organization or contractor and is focused on the development and implementation of the system. An SMP is written by the Program Management Office (PMO) and describes how the software aspects of the acquisition will be managed, what activities are to be performed (including oversight), and how they are performed. Unlike the SDP, the SMP addresses the entire lifecycle of the system. The intent of the SMP is to get program offices to think about the software side of the acquisition much earlier in the acquisition life cycle and not leave those decisions until much later in the overall system acquisition.

Five Key Pillars of Software Assurance Critical to Software Acquisition Management

While the software development processes and practices may vary greatly among programs, our team of researchers in cybersecurity in acquisition and system and software engineering have identified five foundational capabilities that must be well established for effective program results (Woody, 2026). These capabilities (i.e., pillars) are Software Requirements, Software Supply Chain Risk Management, Software Quality, System Integration, and Software Metrics. How each of these capabilities is addressed for a specific program varies greatly, but each one is critical to effective software management. An SMP should include the processes, practices, and responsible resources that will implement each of these capabilities. Inconsistencies, mistakes, and disconnects in these five areas can lead to poor software management decisions and gaps in effective software risk management for both organic software development and software from the supply chain.

These gaps can include improper assumptions about how protections will perform when communications and data sharing occur among the various software and hardware components of the system, including interactions with the hardware platforms that execute the software. There are also gaps in the control of the contents of data packets and other information exchanges that can facilitate attackers. Further misunderstandings occur when designs are created for hardware, but the resulting function is handled by software that may have inherent vulnerabilities. Complex software integrations are assembled from existing components and language libraries that are reused to perform similar actions but may not be built to specification. Also, reused components can include unexpected capabilities and attack surfaces that permit unacceptable behaviors that attackers can leverage. We next consider each pillar and its contribution to software management.

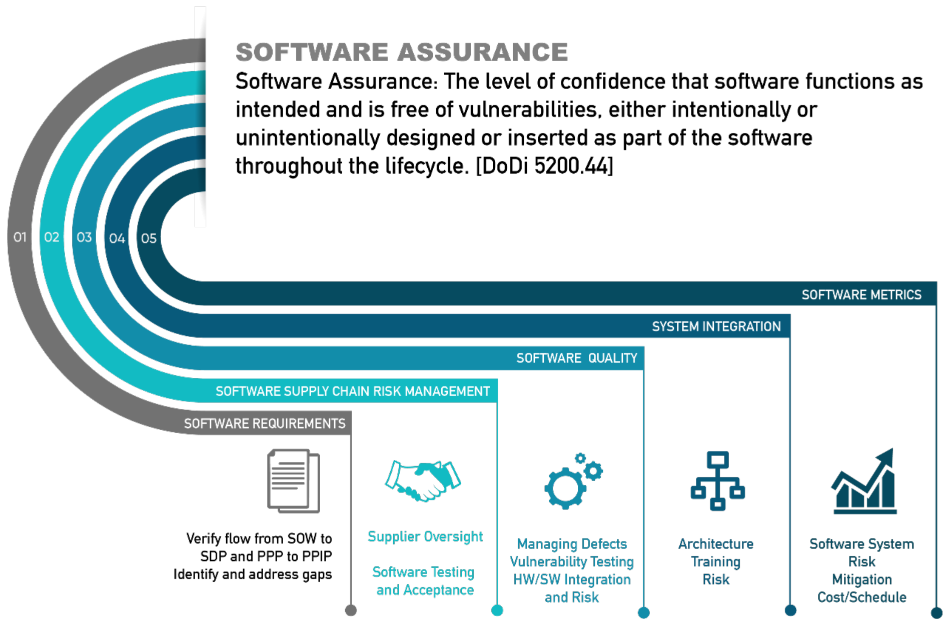

Figure 31. The Five Pillars of Software Assurance Critical to Software Acquisition Management

The first pillar is Software Requirements. This pillar defines the required functionality assigned by the system design to the software/firmware components of the system to perform their allocated capabilities within acceptable constraints and quality to support the system mission. Software developers translate functional requirements into software requirements that feed a software pipeline. The constraints and controls that are placed at the system level should be translated into potential misuse and abuse cases that bound the software and establish what the software will not do and what is not wanted.

Typically, the allocated functional requirements are complete, since they are part of the executable mission that the software must address. However, specific qualities (e.g., performance, reliability, safety, and security) are too frequently incomplete. Decisions made by engineers trained in hardware without software expertise make incorrect assumptions that can severely impact the resulting software. In one program, an assumption was made that the software had no reliability issues because it did not wear out like hardware. This assumption caused the program to ignore software in its safety criticality analysis, which resulted in an insufficient failure mode analysis. Steps for verifying software requirements to ensure they are complete, accurate for the mission, and verifiable need to be part of the SMP.

In addition to the functional requirements levied on a system, there should be other requirements for the development of the system that focus on avoiding system and software vulnerabilities (e.g., the use of secure coding standards, fuzz testing, and penetration testing). Insight into monitoring those additional requirements throughout the acquisition and how they will be verified needs to be a component of the SMP.

Requirements relevant to software can be extracted from the Statement of Work (SOW) and mapped to the SDP to confirm contractor planning, but the SMP should describe the monitoring to ensure these requirements are delivered as planned. Additional software-related cybersecurity requirements are identified in the Program Protection Plan (PPP) and mapped to the contractor-delivered Program Protection Implementation Plan (PPIP), and the SMP should describe how the acquisition will be monitored to ensure effective implementation.

The second pillar is Software Supply Chain Risk Management. This pillar governs and monitors the people, processes, and technology needed to acquire, manage, and sustain third-party software/firmware to ensure it meets system requirements and mission needs throughout its implementation. Using third-party software can offer significant advantages in cost and schedule, but it can also create cybersecurity and life cycle dependency risks that must be managed. When software is acquired, there is an ongoing relationship with the owner of this intellectual property. Managing this vendor relationship needs to be integrated with the system life cycle so that bugs, new features, and vulnerabilities will be addressed by the software owner and transferred to the acquirer securely. It is critical to ensure that the third party can address problems, deliver a software product that shows attention to issues, and use secure transport mechanisms so that the software is not compromised in transit (Ellison et al., 2010). Open source software (OSS) is estimated to be used by 90% of organizations (Gehring, 2022). For OSS, access to the source code is provided, but fixes must usually be posted back to the source provider as required by the license agreement, which provides visibility of these fixes to others who have also installed the code.

Our research has found that decisions about software and its sources are influenced throughout the life cycle in program management, engineering, supplier dependency management, technology support functions, independent assessment and compliance, and process management (Alberts et al., 2023). These processes, practices, and responsible resources for all aspects of software supply chain risk management need to be clearly defined in a SMP to ensure completeness, effective integration, and consistency.

The third pillar is Software Quality. This pillar builds confidence that the people, processes, and practices used to create and acquire software are delivering needed mission results within acceptable levels of rigor and risk. Too often, quality reviews focus only on how well the software development process is performed and do not look at the actual quality of the product that the process is delivering. Higher quality code has been shown to have fewer defects and fewer vulnerabilities, which increases its ability to perform effectively in high-assurance environments (Woody et al., 2014). How quality will be evaluated and who will be responsible for its monitoring should be clearly described in the SMP.

The fourth pillar is System Integration. IEEE defines integration as the process of combining software components, hardware components, or both into an overall system (Institute of Electrical and Electronics Engineers, 2017). Software does not function in isolation, so analyzing and testing it in isolation is only a piece of the puzzle that must be solved to make it functional with sufficient assurance within a system. The process steps and responsibilities for how the software is integrated into the system and verified within that context are too often ad hoc. This provides opportunities for tampering and introducing vulnerabilities into the system because of its integration with other software, firmware, and hardware. An effective integration process should be part of the SMP for both new software and third-party software.

The fifth, and final, pillar is Software Metrics. This pillar identifies and analyzes metrics to demonstrate effective life cycle software assurance to meet mission needs. At the product level, there can be a diversity of metrics that correspond to a particular quality attribute of the system under development and process attributes that relate to the conduct of the engineering activity. These metrics provide objective data that can be used to evaluate the progress, process, and quality of the delivered software. A well-structured SMP will leverage available metrics to monitor acquisition progress and provide early indications of potential problems. As part of our research, we assembled a range of metrics that are typically available from software development activities that can be leveraged in building the specific approach used for software management (Woody et al., 2019).

The selected metrics, which are appropriate to the mission and development approach, are used by the program, and will be analyzed and used for monitoring software, should be clearly described in the SMP. A key aspect to consider is that monitoring must address the full life cycle. Trending will be important in determining whether the qualities the metrics are supporting are improving. The data collected in the early steps of the life cycle (i.e., requirements, design, and architecture phases) can be predictive of how the software will perform, and information collected in later life cycle steps should verify the results (see ***Figure 32***).

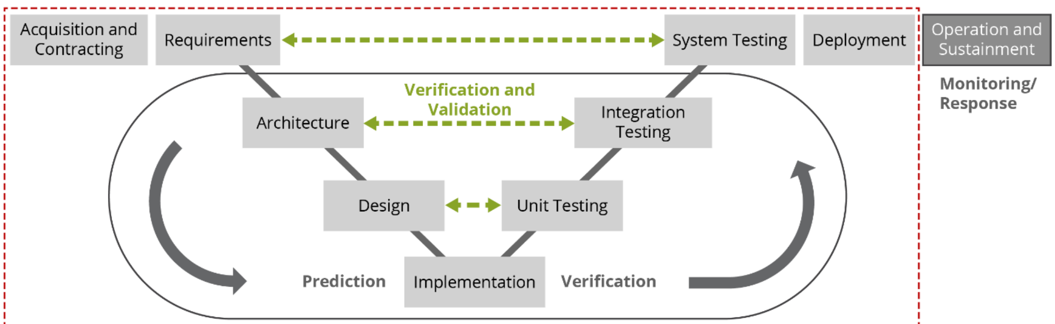

Figure 32. Life Cycle Considerations for Metrics

Defining Threats Critical to Mission Protection

In addition to effective processes and practices for software creation, integration, and monitoring, determining the threats that represent critical risks to mission success need to be analyzed to determine the level of risk that software within the system must address. There are many approaches to threat identification and management. It is critical that the steps for identifying and addressing software threats that will impact the mission be included in the SMP. There are many threat-modeling approaches available. However, the program must ensure that software threats, including those that arise from the software supply chain, are considered. The volume of threats can be overwhelming, and prioritization that balances budget and scheduling constraints with potential impacts of the threats must be considered. The Security Engineering Risk Analysis (SERA) approach (Alberts et al., 2014) can deliver this capability. The SMP will need to include which steps to take to identify and prioritize software threats, handle mitigations, and perform ongoing monitoring to identify when the risk profile for software changes. Software requirements must reflect the mitigations selected for high-priority threats, and enhanced testing should verify that planned mitigations are in place and functioning as expected.

Software Management Plan

In today's technology-driven environments, the importance of software to the proper functioning of a system cannot be overstated. Every acquirer must ensure that the software, which is a key part of the acquired system, is effectively developed and supported to meet mission requirements. This is a complex process. Many acquisitions further increase this process's complexity by spreading the responsibility for software across many parts of the organization. All groups addressing software must collaborate to ensure that they all apply consistent and effective software management. Without effective planning, a well-engineered system that is delivered with software that meets needs consistently and continually cannot be assumed. This planning must describe the processes and practices to be applied across the software within the system's acquisition, which includes the software used in the pipelines to build the software products as well as the third-party selection and the integration process used.

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.

References

Alberts, C., Bandor, M., Wallen, C., & Woody, D. (2023, October 2). *Acquisition security framework (ASF): Managing systems cybersecurity risk (expanded set of practices)* (Technical Note CMU/SEI-2023-TN-004). https://doi.org/10.1184/R1/24128475

Alberts, C., Woody, D., & Dorofee, A. (2014, December 4). *Introduction to the security engineering risk analysis (SERA) framework* (Technical Note CMU/SEI-2014-TN-025). https://doi.org/10.1184/R1/6574856.v1

DoD. (2020, October 2). *Operation of the software acquisition pathway* (DoD Instruction 5000.87). https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500087p.PDF

DoD. (2022, May 13). *Department of Defense standard practice, work breakdown structures for defense materiel items* (MIL-STD-881F). https://quicksearch.dla.mil/qsDocDetails.aspx?ident_number=36026

DoD. (2024, October). *DoD enterprise DevSecOps fundamentals*, Version 2.5. DoD Software Modernization Senior Steering Group. https://dodcio.defense.gov/Portals/0/Documents/Library/DoD%20Enterprise%20DevSecOps%20Fundamentals%20v2.5.pdf

Department of War Under Secretary of War Research and Engineering. (2025, October 28). *Program protection plan outline & guidance*, Version 2.0. Department of War. https://aaf.dau.edu/storage/2025/11/PPP-OG-v2.0.pdf

Ellison, R., Goodenough, J., Weinstock, C., & Woody, D. (2010, May 1). *Evaluating and mitigating software supply chain security risks* (Technical Note CMU/SEI-2010-TN-016). https://doi.org/10.1184/R1/6573497.v1

Gehring, W. (2022). *The state of open source software*. Octoverse. https://octoverse.github.com/2022/

Institute of Electrical and Electronics Engineers (2017, August 28). *ISO/IEC/IEEE International Standard - Systems and software engineering—Vocabulary (*ISO/IEC/IEEE 24765:2017(E)). https://doi.org/10.1109/IEEESTD.2017.8016712

Wallen, C., Alberts, C., Bandor, M., & Woody, C. (2023, June 14). *Software bill of materials framework: Leveraging SBOMs for risk* reduction. https://www.sei.cmu.edu/documents/5364/Leveraging_SBOM_for_Risk_Reduction.pdf

Woody, D., Alberts, C., Bandor, M., & Chick, T. (2026, March 4). *The five pillars of software assurance in system acquisition*. https://doi.org/10.58012/r8q1-zp76

Woody, C. & Bandor, M. (2026, March 11). *Acquisition oversight for software assurance* [Webinar]. Software Engineering Institute (SEI). https://www.sei.cmu.edu/library/acquisition-oversight-for-software-assurance/

Woody, C., Ellison, R., & Nichols, B. (2014, December 22). *Predicting software assurance using quality and reliability measures* (Technical Note CMU/SEI-2014-TN-026). https://doi.org/10.1184/R1/6582113.v1

Woody, D., Ellison, R., & Ryan, C. (2019). *Exploring the use of metrics for software assurance* (Technical Note CMU/SEI-2018-TN-004). https://doi.org/10.1184/R1/12366842.v1

PANEL 23. ACCELERATING GUIDED MUNITION PRODUCTION THROUGH MODULARITY AND SUSTAINED COMPETITION

Thursday, May 7, 2026, 1615 – 1730 ET (1315 - 1430 PT)

Panel Summary:

The return of industrial-scale warfare necessitates new technical and business approaches to acquisition. This panel explores how Modular Open System Approaches (MOSA) and Government Reference Architectures (GRA) are dismantling traditional barriers to entry and enabling a more resilient defense industrial base. By transitioning from single-prime reliance to the direct acquisition of modular components, these research initiatives provide a strategic framework for accelerating production agility—while rigorously examining the cost implications of a decentralized, competitive supply chain and guiding successful implementation.

Chair: Jullian "Dean" Revell, Lead Architect, Air Force Weapons Government Reference Architecture (GRA), Air Force Material Command

Panel Presenters:

A Review of DoW Systems Engineering Modernization Activity – *Dr. Kelly Alexander, Chief Engineer, System Innovation and Nicholas LeGrande, Systems Engineering Modernization Director, OUSW R&E SE&A*

Architecting Affordable Mass – *Gregory Sanders, Joint Production Accelerator Cell / OUSD(A&S)*

Contracting, Acquisitions, and Costing Considerations for Realizing the Promise of MOSA – *Elizabeth Hastings Roer, RAND*

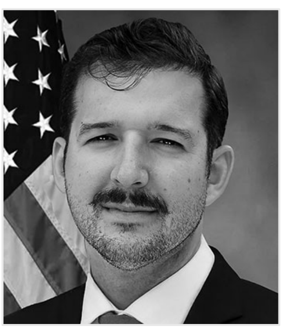

Jullian "Dean" Revell—serves as the Lead Architect for the Air Force Weapons Government Reference Architecture (GRA), a product of the Digital Acquisition and Sustainment Office of the Armament Directorate (AFMC/AFLCMC/EBRD). He leads collaborative efforts with industry and government partners across the DoD to ensure the Weapons GRA and associated products reflect best practices and comply with DoD regulations and standards. The Weapons GRA enables weapon acquisition programs to implement Model-Based Engineering practices and achieve the Air Force's goals of Digital Materiel Management.

Dean has seven years of experience as a civil servant, specializing in systems modeling with SysML, systems acquisition, and test and evaluation. Prior to his current role, Dean was integral to the development of technical requirements and contract language for the Stand-in Attack Weapon (SiAW), a MOSA-compliant weapon system with a development contract valued at $705m. He served as the technical lead for Systems Agility and Architecture as part of proposal evaluation and led MBSE efforts for SiAW post-award. Dean holds a Bachelor of Science in Electrical Engineering and a Master of Science in Systems Engineering from Florida State University.

A Review of DoW Systems Engineering Modernization Activity

Dr. Kelly Alexander—Chief Systems Engineer | System Innovation, Contractor Support to OUSW(R&E) SE&A; kelly.d.alexander12.ctr@mail.mil

Ms. Monique Ofori—Systems Engineering Manager | SAIC, Contractor Support to OUSW(R&E) SE&A monique.f.ofori.ctr@mail.mil

Mr. Nicholas LeGrande—Systems Engineering Modernization Director, OUSW R&E SE&A; nicholas.j.legrand.civ@mail.mil]

Abstract

This paper provides a summary of Systems Engineering (SE) Modernization (SE MOD) published material over the past five years. The SE MOD effort began in 2021 with the intent of understanding the delay and accelerating the implementation of the digital transformation. The Office of the Under Secretary of War for Research and Engineering (OUSW[R&E]) and the Systems Engineering Research Center (SERC) collaborated on this effort that included workshops, one-on-one interviews, and SE surveys across government, industry, and academia. Each of the published reports or papers are summarized below along with a link to access the full report. (Note: Each of the summaries includes references to tables or cited material that can be found in the published report/paper).

SERC 1051 Program Managers Guide to Digital and Agile Systems Engineering Process Transformation

McDermott, T., & Benjamin, W. (2022). *Program managers guide to digital and agile systems engineering process transformation* (Report No. SERC-2022-TR-009). Systems Engineering Research Center. https://www.cto.mil/wp-content/uploads/2023/06/SERC-WRT-1051-2023.pdf

The Systems Engineering Research Center (SERC) conducted a sustained series of research tasks leading to codification of a framework and lessons learned for adoption of Digital Engineering (DE) and Model-Based Systems Engineering (MBSE). DE and MBSE are separate but jointly evolving strategies. DE is defined as "an integrated digital approach that uses authoritative sources of system data and models as a continuum across disciplines to support life-cycle activities from concept through disposal." MBSE is defined as "the formalized application of modeling to support system requirements, design, analysis, verification and validation activities beginning in the conceptual design phase and continuing throughout development and later life cycle phases." Successful adoption will thus be characterized by two general characteristics of the organization: *its ability to conduct systems engineering and related systems modeling, and its effectiveness at digital transformation*.

The SERC research on DE/MBSE adoption found many factors that must be addressed for organizations to achieve this transformation. Table 1 organizes the 12 most prominent factors in our research, organized across three categories: Organizational Design, Organizational Enablers/Barriers, and Organizational Change Management.

The team first conducted an enterprise analysis using interviews to identify what might result from future adoption of DE/MBSE (Section 2.1). Following that, the team used survey and literature reviews to broadly categorize DE/MBSE benefits and adoption factors (Section 2.2).

That study led to a causal analysis that selected primary measurable benefits and adoption factors, resulting in the model in Table 1 (Section 2.3). Finally, additional literature review and interviews were used to conduct an initial validation of Table 1 (Section 2.4).

Systems Engineering Modernization Policy, Practice, and Workforce Roadmaps Lessons Learned and Adoption Framework

McDermott, T., Mesmer, B., & Ergin, N. (2023). *Systems engineering modernization policy, practice, and workforce roadmaps* (Report No. SERC-2023-TR-002). Systems Engineering Research Center. https://www.cto.mil/wp-content/uploads/2023/06/SERC-WRT-1058-2023.pdf

The SERC has analyzed through various surveys, interviews, and literature review the detailed benefits, enablers, barriers, change strategies, and lessons learned that are related to DE/MBSE adoption. As the SERC is a U.S. Department of Defense (DoD) funded research center, the central focus has been on the defense acquisition system and its related industrial base; however, the interviews, surveys, and literature analyses have been conducted broadly across government agencies, industry, and academia. This research conceptualized the changes to workforce and culture necessary to achieve the strategy. A number of qualitative characteristics of DE/MBSE adoption emerged from this study, which are listed in Table 2.

Table 2. Qualitative Statements of Organizational Adoption

- ✓ **Start with SE**: Good SE will enable success with MBSE; MBSE itself will not create value.
- ✓ **Multidisciplinary value:** Systems engineers work across disciplines; DE and MBSE must create added value across disciplines.
- ✓ **Digital literacy**: Experienced SEs may not be comfortable with DE; younger engineers may bring a digital culture with them. Create knowledge transfer between them.
- ✓ **Systems knowledge**: is a unique value of DE/MBSE. A good system model can create/maintain systems knowledge to improve awareness of other disciplinary engineers.
- ✓ **Model quality**: Models must demonstrably improve system understanding and decisions.
- ✓ **Model abstractions**: commensurate to the roles and uses of DE/MBSE. Models must communicate decisions at all levels.
- ✓ **Digital infrastructure**: will make finding and using data more efficient; everything will be on the desktop when needed.
- ✓ **Training**: must be commensurate to the roles and uses of DE/MBSE.
- ✓ **Leadership**: must clearly provide strategy and intent, investment and resources, and the messaging associated with the value of organizational change.
- ✓ **Communication and messaging**: must clearly articulate value to the workforce (benefits) and maintain awareness.
- ✓ **Continuous assessment**: of enterprise capabilities and maturity.
- ✓ **Automation**: of time-intensive tasks will improve value and adoption.
- ✓ **Digital collaboration platforms**: will improve cross-program and cross-disciplinary interaction.

The Supra-System Model

McDermott, T., Alexander, K., & Wallace, R. (2023). The supra-system model. *INSIGHT*, *26*(2), 15–21.

Abstract: This article presents an initial set of concepts resulting from research by the Office of the Under Secretary of Defense for Research and Engineering (OUSD[R&E]) and the

Systems Engineering Research Center (SERC) under an initiative called "Systems Engineering Modernization" (SEMOD). This article discusses the "Supra-system Model," which evolved as a different view of SE life-cycle activities across the entire life of an engineered system. This view promotes SE as a continuous process that is (1) iterative across the full life of a system and (2) managed through a digital transformation centered on data and models. This article also discusses the value of "shared and authoritatively managed data and models" in the life cycle of future systems. These together present a modernized view of SE where "seamless and efficient transfer of data and models" will support practices that are "more agile and responsive to changing stakeholder needs."

An Integration Framework for Digital Transformation of DoD Systems Engineering and Acquisition

Alexander, K., & McDermott, T. (2023, May 10–11). *An integration framework for digital transformation of DoD systems engineering and acquisition* [Paper presentation]. Twentieth Annual Acquisition Research Symposium, Monterey, CA, United States. https://dair.nps.edu/bitstream/123456789/4832/1/SYM-AM-23-063.pdf

Abstract: This paper presents the results of a Systems Engineering Research Center (SERC) research task entitled "Program Managers Guide to Digital and Agile Systems Engineering Process Transformation." This research task supports a larger set of DoD activities being led by the OUSD(R&E) under the term "Systems Engineering Modernization" (SEMOD). The motivation for SEMOD stems from the need to integrate across independent guidance provided down to the DoD SE and acquisition communities related to Digital Engineering, Modular Open Systems Approach, Mission Engineering, and Software Engineering/Agile/DevOps across the multiple pathways of the Adaptive Acquisition Framework. The SERC/government research team found there is a lack of an integrated approach to implementation of SE Focus Areas that is creating a delay in full implementation of the Digital Transformation, which is necessary to ensure the relevant guidance, skills, and training are available to deliver a robust, disciplined approach to weapon systems acquisition.

A Modular Open Systems Approach (MOSA) to Enable Technology Transition

Alexander, K., Ofori, M., & Geier, N. (2025, May 7). *A modular open systems approach (MOSA) to enable technology transition* [Paper presentation]. Twenty-Second Annual Acquisition Research Symposium and Innovation Summit, Monterey, CA, United States. https://dair.nps.edu/handle/123456789/5395

Abstract: Technology Transition is referred to as the "valley of death" due to commonly experienced lack of successful transition to the next phase of system development. The high risk of technology development can cause a delay or cancellation that can be mitigated by Modular Open Systems Approaches (MOSA) principles. MOSA enables technology transition by providing a framework for integrating, upgrading, and replacing components with minimal disruption. By addressing these MOSA principles early in the development cycle, technology transition is more predictable and manageable:

1) Modular Architecture: Focuses on modularity offers plug-and-play capability, where system components adhere to defined standards and interfaces. Modularity also supports incremental upgrades, enabling individual modules to be updated or replaced as technology evolves and parallel development for specified (potentially high-risk) components.

2) Interface Management with Consensus-Based Open Standards: (a) Well-Defined Interfaces that rely on widely recognized, consensus-based open standards, ensuring that new technologies from different vendors integrate effectively, reducing development

and integration challenges. (b) Open Standards facilitate component reuse, which reduces integration time and can also reduce life-cycle costs by increased competition.

3) Enabling Environment that promotes Model-Based Systems Engineering (MBSE) tools processes with access to data to enhance interoperability and options in configuration.

Enabling SETR Modernization Across the Department of Defense (DoD)

Alexander, K., Ofori, M., & Geier, N. (2024, April 30). *Enabling systems engineering technical review (SETR) modernization across the Department of Defense (DoD)* [Paper presentation]. Twenty-First Annual Acquisition Research Symposium, Monterey, CA, United States. https://dair.nps.edu/bitstream/123456789/5112/1/SYM-AM-24-047.pdf

The current SE technical processes and the technical management process remain relevant at the macro level; however, there is a lack of guidance regarding the governance of shared data and models and an understanding of agile and continuous data and model development approaches and their impact on the SETR process. Continuation of document-centric methods will unnecessarily delay acquisition program decisions and impact the ability to implement future modifications and technology insertions. In addition, continuation of the current document-centric methods will not allow the integration of ongoing important initiatives, such as MOSA, software modernization, and Agile program development methods. This study will review the current SETR practices and how digital transformation can improve the process, resulting in better and more agile decision making. The intended outcomes/results of the study include guidance for implementing model-based SETR processes, exemplars for governance and oversight of model-based artifacts, recommended approach to Agile and continuous data and model development to support the SETR, model-based SEP/SETR processes that enable agile continuous data and model development and recommended SE digital artifacts usage and management.

Implementing a Modular Open Systems Approach in Department of Defense Programs

Office of the Under Secretary of Defense for Research and Engineering & Office of Systems Engineering and Architecture. (2025, February). *Implementing a modular open systems approach in Department of Defense programs*.https://www.cto.mil/wp-content/uploads/2025/02/MOSA-Implementation-Guide-v1.1_Public-Release.pdf

This guidebook provides the Department of Defense (DoD) community, including military services, civilians, and DoD support contractors with information to help ensure programs incorporate a modular open systems approach (MOSA) as part of the defense acquisition program life cycle. The intended audience includes decision makers, program management offices (PMOs), program managers (PMs), and program lead systems engineers responsible for implementing and evaluating MOSA in defense programs.

Although MOSA has long been incorporated into DoD programs, this guidebook describes statute and DoD policy that now impose requirements to use MOSA. This guidebook provides best practices for planning, implementing, and evaluating MOSA, including implementation principles, benefits, challenges, and suggestions based on experience from practitioners in the DoD and industry. It includes recommendations to consider when developing requests for proposals (RFPs) and evaluating planned approaches for implementing MOSA. The appendices provide supplemental information on the background efforts that have helped inform the DoD MOSA community.

Leveraging Modern Systems Engineering Principles to Enhance the Department of War Software Acquisition Pathway

Leveraging Modern Systems Engineering Principles to Enhance the Department of War

Software Acquisition Pathway: https://www.cto.mil/sea/mosa/

The Department of War (DoW) is at a pivotal moment in modernizing its software (SW) acquisition practices to meet the demands of an evolving threat landscape and the rapid speed of technology advancement. In March 2025, the Secretary of War (SECWAR) directed that software acquisition emphasize speed, flexibility, digital integration, and the rapid infusion of commercial innovation via the Software Acquisition pathway. The SECWAR released additional guidance regarding acquisition transformation and reforming the requirements generation process. In parallel, the Office of the Under Secretary of War for Research and Engineering (OUSW[R&E]) is leading Systems Engineering Modernization (SEMOD) to implement digital transformation across technical reviews and systems engineering (SE) process flows to support the acquisition decision process. Each of the SEMOD principles align with the SECWAR transformation initiatives by providing the methods to accelerate capabilities to the warfighter while also ensuring the ability to allow future upgrades. This paper provides considerations and metrics for integrating rigorous, modern SE principles into the Software Acquisition pathway, enabling rapid, iterative delivery of software capabilities for timely, data-driven acquisition decisions.

To successfully execute the Software Acquisition pathway, programs should adopt modern SE practices in both the Planning and Execution phases that align with the department's digital transformation initiatives. The OUSW(R&E)'s SE MOD effort provides the foundation for this transformation by integrating digital engineering, modular open systems approach (MOSA), mission engineering, and Agile Development, Security and Operations (DevSecOps) practices into program planning and execution. These practices are critical for maintaining authoritative digital models, enabling data-driven decisions and ensuring program teams can respond rapidly to evolving requirements without compromising cost, schedule, or performance. To enhance the execution of the Software Acquisition pathway and deliver timely, resilient, and upgradable capabilities to the warfighter, program managers (PMs) and chief systems engineers (CSEs) should consider these SE MOD principles when applying digital, iterative, and model-centric approaches.

Prepare and Develop the Workforce for a digital environment by adopting a "digital -first mindset."

- ✓ Modernize the Workflow to align with digital practices and tools.
- ✓ Implement MOSA principles to manage architecture that supports future upgrades, modifications and enables integration.
- ✓ Adopt Modern Design Practices (relying on Agile and Lean practices) to apply an iterative approach to design, test and incremental capability release that is informed by continuous user input.

The referenced documents above provide the necessary definitions, insights and recommendations that support a robust understanding of the role SE Modernization, digital transformation and MOSA have on DoW acquisition. The referenced material and additional SE MOD guidance is available at cto.mil.

Architecting Affordable Mass

Gregory Sanders—is a Senior Policy Analyst with the Joint Production Accelerator Cell in the Office of the Under Secretary of War for Acquisition and Sustainment. He also works to surge munitions production through coproduction and other means, with a particular interest in Modular Open System Approaches. Sanders is on an Intergovernmental Personnel Act detail from the Center for Strategic and International Studies, where he is a senior fellow. He holds a master's degree in international studies from the University of Denver, a bachelor's degree in government and politics, and a bachelor's degree in computer science from the University of Maryland. [Gregory.S.Sanders27.civ@mail.mil]

Abstract

The U.S. guided munitions enterprise faces critical production and surge limitations, a significant vulnerability in the context of industrial-scale warfare and strategic competition. While current efforts to ramp production through larger budgets and supply chain investments are underway, these solutions are often slow and capital-intensive. They are further hampered by a consolidated industrial base, high barriers to entry, and boom-and-bust demand cycles that discourage investment.

This paper examines how the synergy between two key initiatives—affordable mass programs and Modular Open System Approaches (MOSA)—can address these challenges by fostering production adaptability. Affordable mass programs leverage flexible requirements and commercial technologies to boost producibility and lower unit costs. In parallel, MOSA uses government-defined open architectures to lower barriers to entry, sustain competition, and speed innovation. This paper draws from interviews regarding pathfinding munitions and adjacent systems to draw lessons on achieving production capacity and surge capability. The paper identifies seven categories of enablers: (1) enhance module competition, (2) promote reuse, (3) modularize in-service, (4) commoditize modules, (5) sustain markets, (6) sustain system competition, and (7) acquire mobilization options. The paper explores how these enablers interact and suggests options to scale the approaches of pathfinder programs.

Executive Summary

The global proliferation of munitions, advances in unmanned systems and their use in conflict, China's emergence as a manufacturing superpower, and the return of industrial scale warfare are all pressing problems for the U.S. guided munitions enterprise. The Department and the defense industrial base (DIB) are limited in their ability to produce and surge guided munitions, especially those with exquisite performance. Executive orders, legislative actions, and departmental decisions have sought to address this problem.

In a panel discussion on this topic, Maj Gen Kunkel drew from his own experience as a fighter pilot to outline what limitations look like to an operator in wartime. In the initial weeks of the war in Kosovo, the Air Force had top-of-the-line missiles that could hit a target the size of the projector in the room, but after a few weeks those were expended, and the targeting accuracy was decreased to a target the size of the tall and wide screen. A few weeks later they were reliant on stockpiles of dumb bombs that may hit a target the size of the room. This challenge is even more dangerous in a great power conflict such as the war in Ukraine, where defending vital targets is increasingly difficult, let alone to deterring a war with China.

Elected leaders and the Department are already pursuing the munition production growth via one primary and one supporting avenue, respectively:

- **Increasing the demand signal** through budgets, block buys, and the use of multiyear procurement (MYP).

- **Investigating in supply chains and intervening** to address limiting factors and bolster the munitions industrial base (IB), including both private and government-owned facilities.

These steps are already boosting production capacity. However, that process is capital-intensive and time-consuming. This inherent difficulty is magnified by three additional problems:

- **High barriers to entry exist** due to inherent technical and qualification difficulty, magnified by incomplete data and a consolidated munitions IB segregated from commercial industry.
- **Production capacity is fragmented by program**, leading to notable nonrecurring engineering, time, and financial costs to crossover, especially for the qualification process.
- **Boom-and-bust cycles and uncertainty** impede investment. The demand for munitions regularly changes, and the industrial economics only reward addressing enduring demand.

Fortunately, the Department is already pursuing two initiatives that, along with supply chain efforts, enable production adaptability. These efforts provide combat capability by accelerating the rate at which the acquisitions system and IB orients and acts on operator needs.

- **Developing affordable mass programs** that leverage flexible requirements, commercial technology, and new manufacturing processes to boost producibility and lower unit costs.
- **Architecting openness and commonality** to lower barriers to entry, sustain competition, and speed innovation through a modular open system approach (MOSA), government reference architectures (GRAs), or government reference designs (GRDs).

This project examines how synergies between how affordable mass and MOSA can address the warfighting imperative of production adaptability. The results are summarized in ***Table 12***.

Table 12. Theorized Affordable Mass and MOSA Pathfinder Enablers of Production Capacity and Surge Capability

Enabler	Description	Production Capacity	Surge Capability	Pathfinding Examples
Enhance Module Competition	Allow open competition via open architectures with severable modules	● Adds flexibility to route around bottlenecks and rewards speed of scaling	◑ Incentivizes producibility but optimizing for price risks losing slack capacity	Enterprise adoption of WOSA and AMS GRA
Promote Reuse	Use technology and items from any supplier across programs and life cycle	◑ Reduces development and test requirements but requires conformity	○ Creates commonality which allows for other benefits but has risks	SOSA focus on scale of adoption; GPS modules
Modularize In-Service	Incrementally open a portion of a system's architecture.	◑ Addresses bottlenecks, but requires leverage, e.g. block upgrade or MYP	◑ Enables pursuit of other benefits, as above	Exquisite munition mission systems and propulsion 2nd sourcing
Sustain a Market	Maintain diverse IB by pooling demand and enabling interchangeability	◑ Leverages commercial IB, small business innovation, and allied IB	● Preserves engineering communities and allows for competition for scaling	FACE and SOSA for software and sensors respectively
Commoditize Modules	Prioritize producibility and lower qualification risk	● Sacrifices other priorities for production ramping	◑ Lowers barriers to entry and eases surging	Energetic subsystems experimentation
Sustain System Competition	Use government reference architectures to sustain or expand prime competition	◑ Avoids prime vendor lock, but challenging for exquisite systems	○ Eases bringing in latent or international capacity, but requires development	Affordable mass missiles and UAS; CCA; OMS/UCI inter-system MOSA
Option to Mobilize	Employ contract-producers to address production constraints	◑ Allows for mobilization, if license or government reference design in place	● Mobilizes flexible manufacturing and partner industrial bases.	Development of energetic subsystem technical data packages; franchise model
Legend: ○ Relationship situational or faces high barriers; ◑ General benefit but notable risk(s) or scope limit(s); ● Clearcut benefits				

Source: 18+ interviews with pathfinding programs and guidebooks (Fookes, 2023; Guertin, 2025; Kendall, 2013; OSE&A, 2025). Notes: The table draws from munition and adjacent systems, standards, and architectures with relevant to MOSA. The key standards and architectures include the Weapon Open System Architecture (WOSA) and Weapon GRA, the Open Mission System/Universal Command and Control Interface (OMS/UCI), the Agile Mission System, the Future Airborne Capability Environment (FACE), the Sensor Open Mission System Architecture (SOSA). For the Army, this project focuses on Army Launched Effects (LE) / Long-Range Precision Munition (LRPM). For the Navy, this project examined Conventional Prompt Strike (CPS) and the nascent Coalition Heterogenous Affordable Offensive Strike (CHAOS). For the Air Force, the project considered the Air Force's Collaborative Combat Aircraft (CCA), Family of Affordable Mass Munition (FAMM) and Extended Range Attack Munition (ERAM), and Dragon Cart. Multi-service and Research and Engineering efforts to expand production of energetics subsystems, including the Low-Cost Cruise Missile (LCCM), also provide important examples at the subsystem level.

Introduction

Advances in unmanned systems and their use in conflict, the return of industrial scale warfare, and growing demands on air and missile defense systems are pressing problems for the U.S. guided munitions enterprise. This challenge is reinforced by China's rise as a manufacturing superpower as the U.S. economy shifts to services as part of globalization. Current U.S. guided missiles have exquisite capability but face challenges that limit the defense industrial base's (DIB's) production capacity and surge capability. Executive orders, legislative action, and departmental decisions have invested time and money in pursuit of addressing this problem.

In a panel discussion on this topic, Maj Gen Kunkel drew from his own experience as a fighter pilot to outline an operator's wartime perspective on these limitations. In the initial weeks of the war in Kosovo, the U.S. Air Force had top-of-the-line missiles that could hit a target the size of the room's digital projector. After a few weeks those were expended, and the available munitions could instead hit a target the size of the tall and wide screen. A few weeks later, operators were reliant on stockpiles of dumb bombs that may hit a target the size of the room. This challenge is even more dangerous in recent industrial-scale conflicts, let alone to deterring a war with China.

Elected leaders and the Department are already pursuing the production aspects of the munition production challenge via one primary and one supporting avenue, respectively:

- **Boosting the demand signal** via budgets, handshake deals, and multiyear contracting.
- **Investigating supply chains and intervening** to address limiting factors and bolster the munitions industrial base (IB), including both private and government-owned facilities.

These steps are already boosting production capacity. However, that process is capital-intensive, requires time-consuming facilitation efforts, and faces difficult workforce constraints. This inherent difficulty is magnified by three additional problems:

- **High barriers to entry exist** due to inherent technical and qualification difficulty, incomplete technical data packages, and a consolidated DIB segregated from commercial industry.
- **Production capacity is fragmented by program**, leading to notable nonrecurring engineering, time, and financial costs to crossover, especially for the qualification process.
- **Boom-and-bust cycles** impede investments. The demand for munitions regularly changes, and the industrial economics only reward addressing enduring demand.

This paper explores the production acceleration and returns on investment that could be achieved by a Modular Open System Approach (MOSA). MOSA's established benefits have the potential to address these challenges. However, policy-makers lack good metrics and models to weigh the relevance of MOSA to their production problems.

This paper further explores how developing affordable mass programs that leverage flexible requirements, commercial technologies, and new manufacturing processes boost producibility and lower unit costs. Affordable pathfinders used two more approaches of interest:

- **Sustaining System Competition** by employing government reference architectures (GRAs) to retain multiple primes through production to allow for future competitive on-ramps.

- **Acquire Mobilization Options** via expanding production using franchising / licensing options or government-held technical data packages to address latent industrial capacity.

Key Concepts

MOSA: Programs employing MOSA have several modules that are internally highly cohesive while being loosely coupled with one another through open interfaces. For example, a guided missile may have a solid rocket motor as one module, the warhead as another, and the guidance system as a single subsystem or multiple components. MOSA draws on commercial principles seen in Android phones, universal serial bus (USB) devices, and International Business Machines (IBM) personal computers. The use of a MOSA, mandated for new programs of record to the extent practicable by 10 U.S.C. 4401–4403 and reinforced by tri-service memos, involves technical and business approaches that "enable incrementable development and enhanced competition, innovation, and interoperability" (Del Toro et al., 2024, p. 2).[1] MOSA standards are collected in GRAs developed and shared with partners in industry. MOSA is easiest to apply to a developmental program and has benefits throughout the life cycle but is traditionally not focused on production capacity or surge capability.

Affordable Mass: The concept of affordable mass originates from the idea that the Department should invest in a high–low mix, which is a combination of low-cost systems and exquisite systems that complement each other.[2] Affordable mass systems need flexible requirements that accept niche or lower-level performance to enable greater adoption of commercial technology and advanced manufacturing approaches. Precision-guided munitions and attritable uncrewed systems are well-suited to affordable mass because human pilots are not at direct risk. Finally, affordable mass systems rely on software to allow for fast and scalable adaptation. The name captures two key metrics of success: *affordable* refers to lower unit cost while *mass* implies a mix of production capacity or surge capability. Affordable mass programs must avoid or overcome factors that slow the scaling of exquisite systems.

Production Capacity and Surge Capability: Production capacity refers to the quantity and timeline of goods that the DIB produces under current funding and prioritization. Increasing production capacity involves slow and expensive investments. Surge capability refers to the ability to rapidly increase production. This may be accomplished through employing idle DIB capacity from current producers, including facilities and tools that are set aside ("mothballed") for emergency use. Alternatively, this could involve providing a technical data package (TDP) to a new producer with flexible manufacturing capabilities in private industry or the organic industrial base. This would require a franchising or license agreement or government-owned TDPs or interventions (Cook, 2023).

Research Issue and Method

This paper is concerned with the combat capability provided by DIB production capacity and surge capability. MOSA guidebooks discuss scalability, obsolescence, and diminished manufacturing sources and material shortages (DMSMS) but not production capacity or surge capability directly (Fookes, 2023, p. 15; Guertin, 2025, pp. 14–16; Kendall, 2013; pp. 25, 48;

---

[1] The tri-service memo charges acquisition officials with implementing five pillars of MOSA: "(1) employing a modular design, (2) designating modular interfaces, (3) leveraging consensus-based open standards, (4) establishing enabling environments, and (5) certifying conformance" (Del Toro et al., 2024, p. 2).
[2] The concept of a high–low mix, focused on aircraft, dates back to reform proposals in the 1970s and 1980s (Carroll, 2008). Recent discussion of high–low mixes and the newer term *affordable mass* tend to emphasize not just weapon cost but also producibility (Gunzinger, 2001; Pettyjohn, 2025).

OSE&A, 2025, pp. 7, 23–24).[3] While MOSA's benefits are relevant to production capacity and surge capability, the Department's focus on production acceleration demands trade-offs under tight time horizons. This paper studies pathfinding MOSA programs to better understand how and to what extent a MOSA approach allows the production of more guided munitions faster. The author conducted 18+ interviews to answer the following research questions:[4]

- How do munitions and adjacent programs that employ MOSA incorporate producibility into their architectural decisions and investments?
- What traditional munitions producibility challenges are addressable through modularity or adoption of commercial standards, and what are the present plans of affordable mass programs to address or bypass these limitations?
- To what extent could architectural design expand DIB production capacity or surge capability by incorporating nontraditional or international producers or modernizing the existing U.S. DIB?

Results

The interviews with pathfinding programs found that MOSA is being experimented with by a mix of exquisite and affordable mass munitions. Production capacity and surge capability was typically a secondary consideration. Nonetheless, both pathfinders and guidebooks agreed on addressing production constraints using consensus-based open standards that allows cross-vendor interoperability, break-out modules for components of concern, and substitution of components with ones from other suppliers (Fookes, 2023, p. 15; OSE&A, 2025, pp. 23–24). The first three enablers draw from MOSA guidebooks:

1. **Enhance Module Competition:** Using severable modules with open architectures.
2. **Promote Reuse:** Reducing testing and development by reusing items across programs.
3. **Modularize In-Service:** Employing incremental MOSA to allow reuse in exquisite systems.

The remaining enablers are drawn from MOSA pathfinders:

1. **Commoditize Modules:** Prioritizing producibility and reducing risks that lengthen qualification for items that are capital intensive to manufacture and lack dual-use markets.
2. **Sustain Markets:** Building competitive and pooled markets for modules and resilience by making it easier for multiple suppliers to achieve a minimum sustainable production rate.
3. **Sustain System Competition:** Retaining competitors and the ability to employ and on-ramp primes via GRAs and affordable mass approaches.
4. **Acquire Mobilization Options:** Ensuring means to scale manufacturing in ways not required by MOSA—for example, through a franchising or manufacturing licensing agreement option or by arranging for a distributable government reference design.

[3] The *DoD Producibility and Manufacturability Engineering Guide* mentions MOSA once, noting that it will be considered as part of the Independent Technical Risk Assessment for Major Defense Acquisition Programs (OEDSE&A, 2024).

[4] See the appendix for the guided conversation questionnaire used in these discussions.

The remainder of this paper dives into these enabler categories to draw findings from pathfinder experimentation.

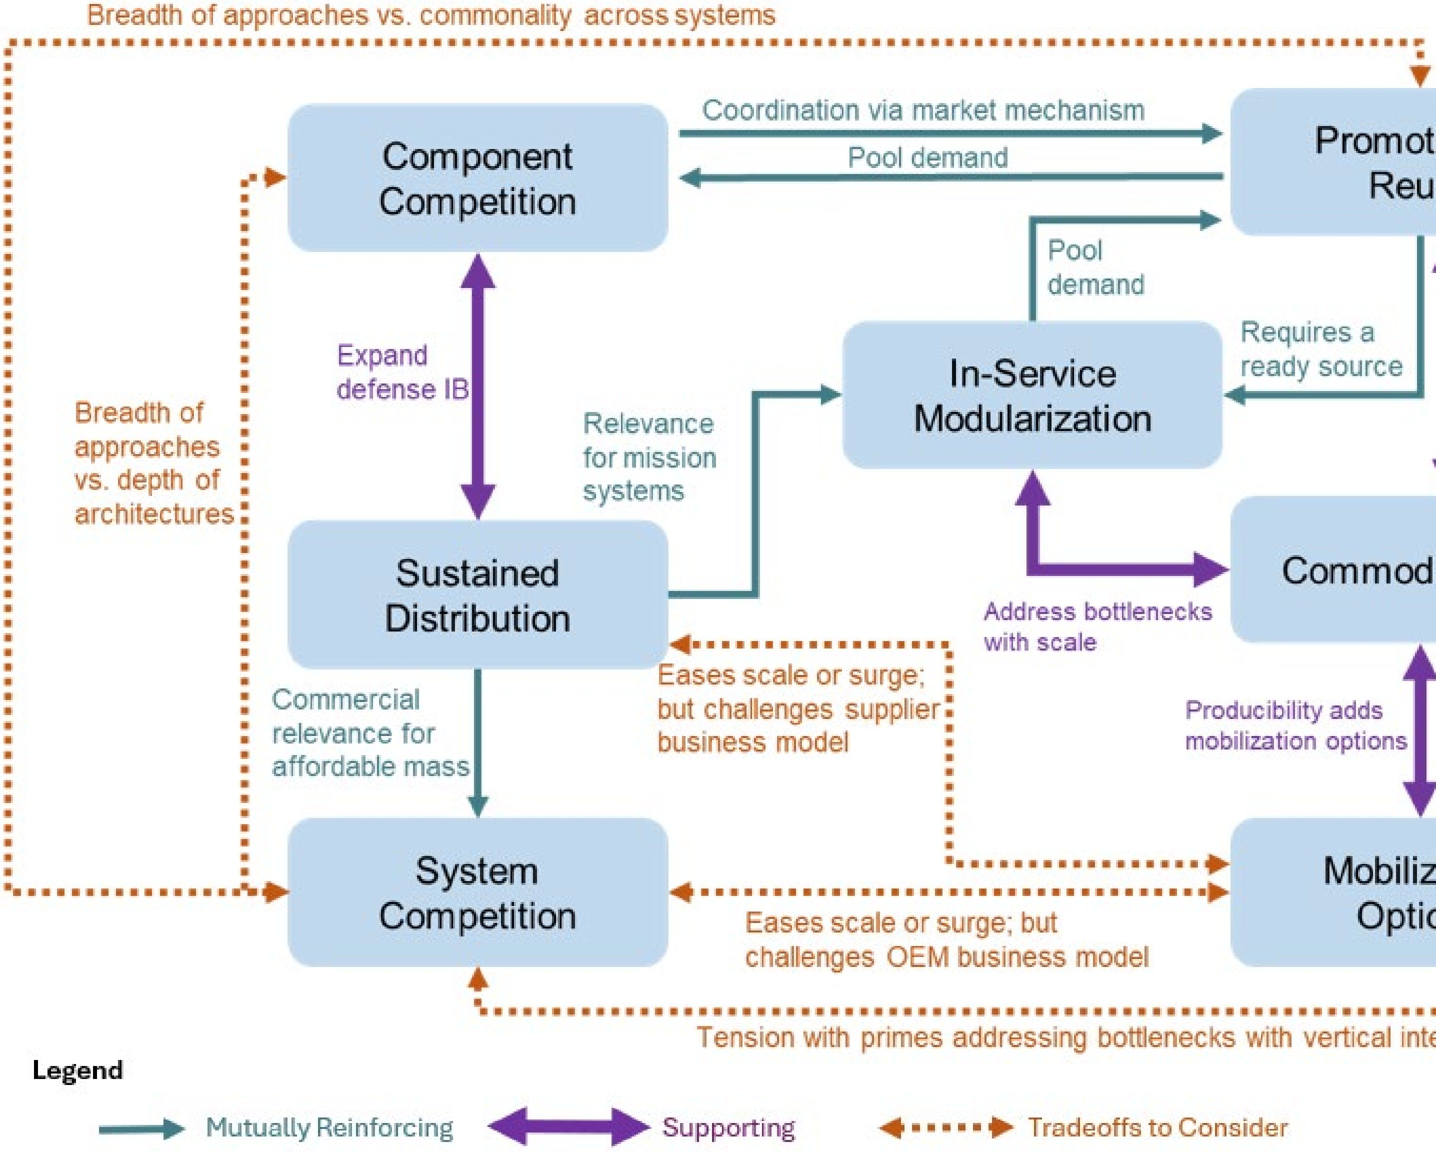

Figure 33 presents a theorized relationship between these tools.

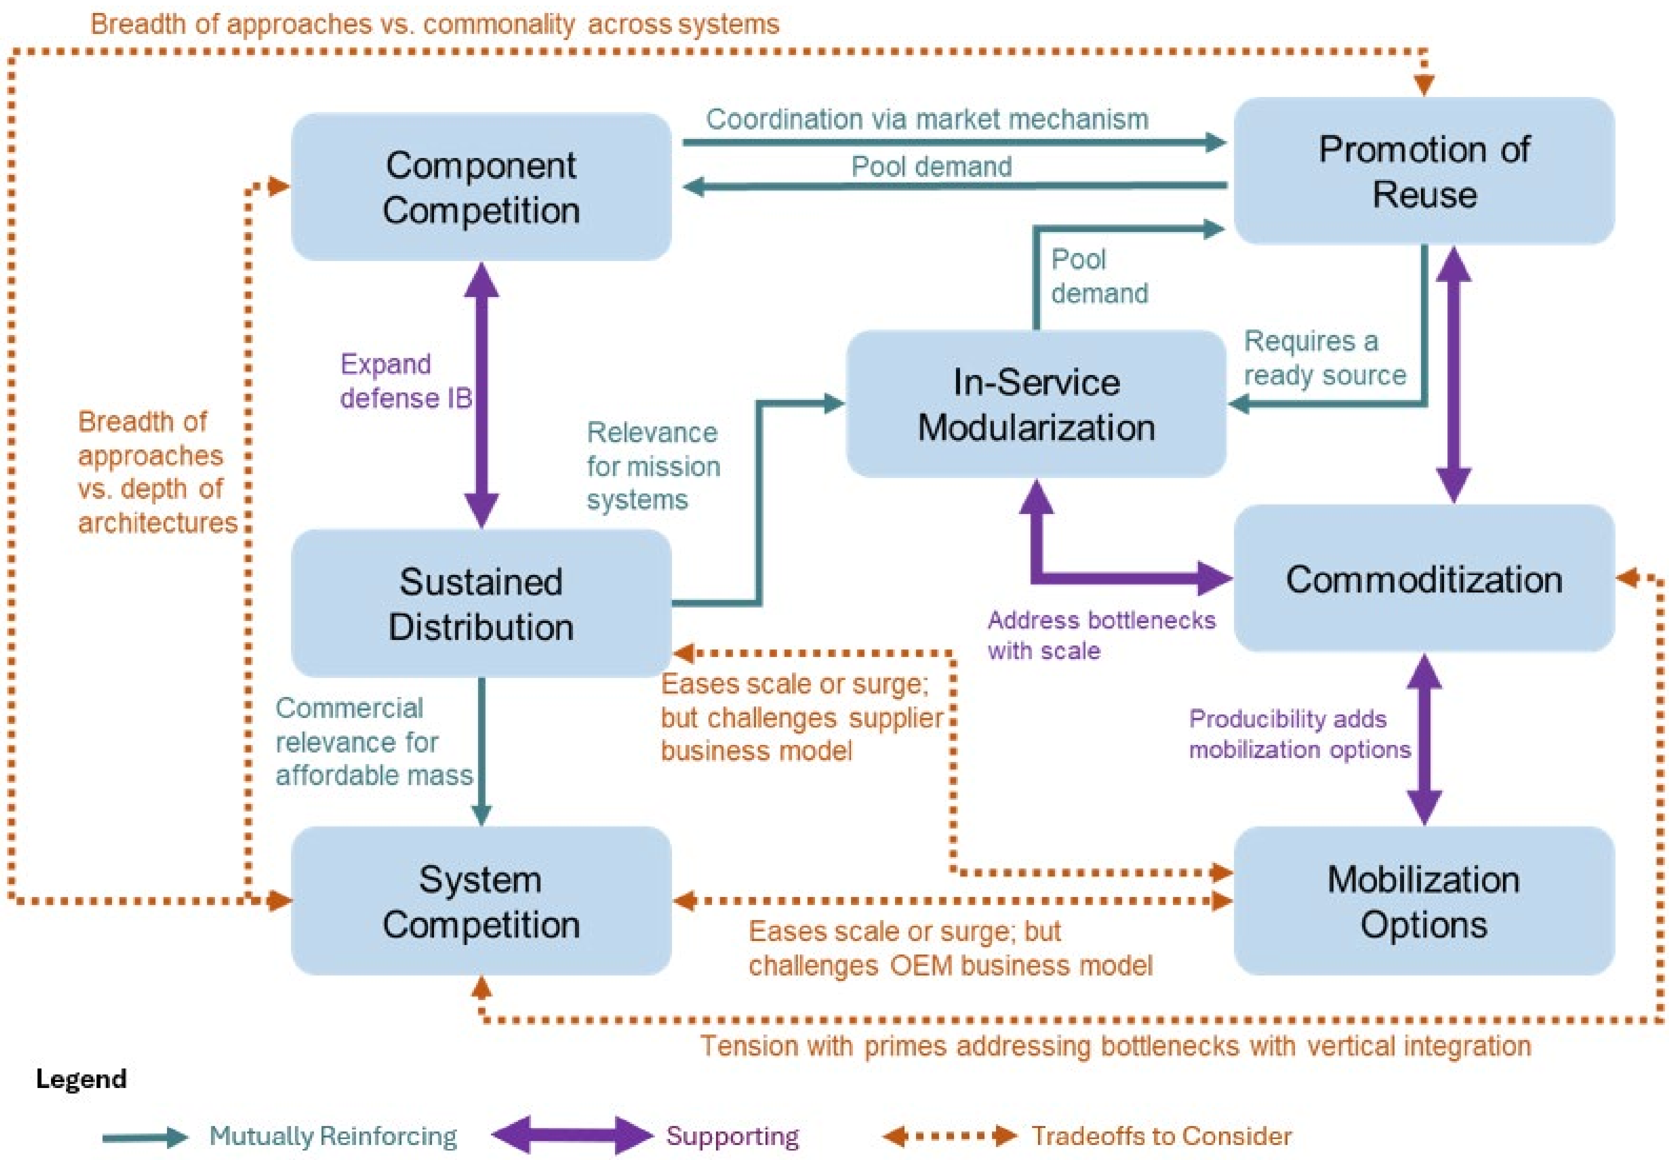

Figure 33. Interaction Between Affordable Mass and MOSA Pathfinder Enablers of Production Capacity and Surge Capability

Enhance Module Competition and Promote Reuse

Component competition is core to MOSA and enabled by "open architectures with severable modules" that may be substituted. A diverse and competitive supplier base "fosters innovation and drives down costs" (OSE&A, 2025, p. 6). The possibility of competition gives the government flexibility if another supplier is available and lowers barriers to the entry of new suppliers. However, even competitive markets experience single points of failure, and qualifying new suppliers will be costly and time-consuming, especially where domestic production does not have a comparative advantage.

Reuse of components between programs or variants is also enabled by an open architecture with severable modules. Widespread adoption of a MOSA standard "promot[es] the reuse of technology, modules, and components from any supplier across the acquisition life cycle, reducing the need for redundant development efforts, reducing test requirements, and leveraging existing investments efficiently" (Guertin, 2025, p. 5). Enterprise mandates and mature standards enable pooled and competitive supplier markets (Fookes, 2023, pp. 22–23). That said, enterprise standards do not guarantee reuse because two programs may implement the same standard in different ways that adds friction to commonality.

The most relevant standard for munitions is WOSA, which was developed out of the Air Force Research Lab with help from an industry consortium starting in 2014. This standard and associated weapons GRA has been adopted by multiple U.S. Air Force and U.S. Navy programs, discussed below, each with its own implementation architecture. The Munition Open Architecture Test and Evaluation Laboratory (MOATEL) maintains the standard and provides expertise and modeling capacity to assist in employing and verifying the standard.

WOSA has wrestled with the trade-offs of promoting component reuse versus seeking widespread adoption. After initially taking a more rigid approach, the architects chose to open the standard up to a wide range of projects and focus on standardizing the logical messaging. This left key hardware and transportation protocol decisions to the vendors and mission space. Timing was an important concern when defining the score of domains under WOSA. The architects used research and simulation to ensure that the dividing lines between modules did not result in requirements that industry would not be able to follow. Architecting for affordable mass munitions is less demanding than for more exquisite systems that deal with larger temperature ranges or faster speeds. As a result, an affordable mass system could mandate commercial ethernet standard connections or physically modularize almost every hardware subcomponent. WOSA is now working to extend this logical model into the physical domain.

In addition to WOSA, three other standards highlighted by the tri-service memo's shortlist were included as pathfinders (Del Toro et al., 2024, p. 1):

- The Agile Mission Suite Government Reference Architecture (AMS GRA) is the U.S. Air Force standard for crewed and autonomous air platforms. AMS GRA deals with large clusters of subsystems. For example, in the Collaborative Combat Aircraft (CCA), there are modules for the air frame, mission systems, and autonomy.
- The Sensor Open System Architecture (SOSA) is a standard for command, control, communication, computers, cyber, intelligence, surveillance, and reconnaissance (C5ISR) payloads. SOSA coordinates U.S. government customers to agree on high volume common form factors and solutions that deliver affordable C5ISR payloads.
- Future Airborne Capability Environment (FACE), like SOSA, emphasizes commonality. FACE decomposes software applications to modules and manages data flows.

The U.S. Army was, at time of interviews, experimenting with their approach to uncrewed systems and missiles. The U.S. Army Launched Effects (LE) is grouping some missile programs within their larger unmanned aerial systems (UAS) family of systems and enterprise architecture approach, including FACE. However, attempts to apply MOSA to small UAS have encountered challenges and may focus on external interfaces (i.e., command and control, and dispenser/launcher) and system competition rather than modules within systems.

Modularize In-Service and Commoditize Long Lead Items

Bringing in a new supplier or manufacturing approach is challenging if that supplier or approach has not already provided equivalent or close components to a different program. The difficulty is magnified for systems that do not have an open interface for the module in question. Most in-service guided munitions have not yet implemented MOSA and tend towards highly intertwined functions. While MOSA is most easily incorporated early in the life cycle, when the government has leverage, this barrier can potentially be overcome by opening a loosely coupled portion of the architecture via incremental MOSA (Fookes, 2023, pp. 13–14; Guertin, 2025, p. 6).

Even in the absence of MOSA, primes will seek out second sources, sometimes in response to inducements, pressure, and/or change orders from the U.S. government. Prime vendors do have an incentive to address production constraints and to maintain surge capability, because producing at the pace of demand would earn revenue faster. However, expanding production capacity and maintaining slack capacity are both expensive, and given the history of munition boom-and-bust cycles, a risk averse company may prefer to maintain a backlog, keep costs low, and win sales over a longer time frame.

Sustained ecosystems drawing on the commercial sector offer the most opportunities for innovation. However, when dealing with military-specific technology to address a critical bottleneck, one program at a time may be the fastest option. This might involve an exquisite capacity, for example Military-Code (M-Code) GPS receivers that can overcome jamming but faced significant delays (Ludwigson, 2024). Alternatively, this may involve a widely produced item with difficult-to-source materials, constrained lower-tier suppliers, like solid rocket motors (Shalom, 2026). Multiple factors add time and cost to the qualification process: certify a new vendor or production line, a new production process, or a new design. The standards are even higher for flight-critical components, where any failure could lead to unsafe operations, or energetics, which involve dangerous chemicals and the risk of explosions.

Incremental adoption of MOSA could help scale the benefits of direct-to-supplier efforts to address a critical bottleneck because GRAs and domain implementation packages can be applied to other programs or vendors.

Commoditization is a pragmatic approach for bottlenecks involving difficult-to-qualify items because it focuses innovation on widespread producibility rather than specialized performance. Affordable mass systems have an edge in this case because more of their performance is defined by software, allowing for more reliance on lower-cost mass-produced hardware. However, even for exquisite systems, a subsystem like a solid rocket motor will have components such as inert motor cases where qualification could be eased through greater use of proven items across programs, albeit still needing to allow for differences in diameter or the like.

Sustain Markets

A diverse and distributed manufacturing ecosystem has considerable advantages for resilience and is an opportunity to build up geographic clusters of producers that advance process knowledge that is key to manufacturing success. Competing producers can uncover new production approaches, mitigate risks to any given facility, and offer more choices when it is time to surge.[46]

The number of competitors that can be economically maintained for a system or module depends on the minimum sustainable rate of production. A competitive and pooled market for modules makes it easier for multiple suppliers to achieve that rate and allows primes to take advantage of a shared supply chain. Sustaining distribution is most straightforward when vendors can draw on commercial markets or allies subsidizing sovereign production, or when small businesses can manufacture the item on their own.

The maximum number of sustainable competitors will depend on the extent of demand and the magnitude of barriers to entry and capital expenses—that is, how many vendors could reach a minimum sustainable rate of production. MOSA expands this number by pooling demand to create larger markets, including dual-use technology, increasing the frequency of competition, and offering on-ramps for new vendors. Pathfinders have made decomposition and disclosure choices that allow access to wider commercial or international markets.

Decomposition refers to how a larger system is broken down into smaller modules, which has functional and physical aspects. FACE's and SOSA's granular modules and adoption of commercial standards make them more accessible to small businesses. Incremental MOSA can likewise target modules of less concern to primes with larger potential markets, such as gyroscopes. The greatest benefit for surge capability comes if there is a robust commercial-off-the-shelf (COTS) or near-COTS market for the module in question—for example, for microelectronics or hardware connectors. Affordable mass systems systematically seek to include commercial standards and COTS, though some of the approaches may scale to exquisite systems willing to accept shorter item lifespans.

Disclosure of standards can also ease the participation of international vendors, in line with the National Defense Strategy.[47] U.S. allies and partners have been increasing their defense spending and purchases of U.S. weapon systems, but electoral and sovereignty concerns mean that access to these markets will often depend on partner industrial participation. In addition, MOSA designs for exportability, which enables the substitution of less sensitive items for export, such as seekers, reducing risks of variant-lock and easing the strategic leveraging of partner industrial bases while protecting technology security.

The CCA has been an exemplar of employing MOSA to better leverage international production capacity and investments to ramp production and build surge capability (Decker, 2025; Sanders & Aldisert, 2024). Standard accessibility to allied and commercial industrial bases has been furthered by the U.S. Air Force making an unclassified version of AMS GRA and the release of WOSA to select partners. The FACE standard has been released to the public, incorporated into British standards, and includes Five Eyes industry in their consortium.

Sustain System Competition

This enabler refers to the ability to sustain and on-ramp multiple competing primes for production. This is enabled by shared GRAs and is most common for affordable mass systems.

---

[46] Manufacturing ecosystems are an area where China has overtaken the United States in some sectors. For example, Shenzhen is a global electronics manufacturing powerhouse (Wang, 2025, pp. 59–68). That said, economies of scale are a powerful force, and even as China has expanded its ballistic missile facilities and production capacities over the last two decades, their number of companies and facilities has consolidated (Jones & Palmer, 2024, pp. 27–28; Wood & Stone, 2021, pp. 2, 19–20).

[47] The National Defense Strategy mandates "leverag[ing] allied and partner production not just to meet our own requirements but also to incentivize them to increase defense spending and help them field additional forces as quickly as possible" (Hegseth, 2026, p. 4).

Promoting reuse via commonality allows for shared supply chains, but greater variation may allow for more primes and approaches to producibility.

The traditional departmental approach is to compete primes against one another to win development contracts using a winner-take-all model for production. Vendors are incentivized to bid aggressively to win contracts and then to make their money back in production and, primarily, sustainment. Purchasing from multiple primes and sustaining multiple supply lines is expensive. Modularity can mitigate these challenges, but the largest multi-award crewed system of this century, the Litoral Combat Ship, experienced significant production challenges and met an early end. Other crewed systems like the Future Long-Range Assault Aircraft were able to leverage MOSA and other transaction authority to sustain competition longer than typical but still chose a single prime.

Today, the uncrewed CCA and affordable mass systems are leveraging MOSA and GRAs to sustain competition into production and, in theory, maintaining the ability to on-ramp additional primes in the future. The GRAs make it easier to maintain relationships with multiple vendors and for them to rapidly iterate the products they are experimenting with into compliance with a given program. In its first tranche, the CCA is a notable success story for AMS GRA, going from contract through development in 14 months. Some affordable mass munition and small UAS programs are maintaining a variety of vendors and intend to carry that forward or allow for on-ramps. MOSA allows systems to go beyond interoperability to interchangeability. Maintaining competitors is easier if operators and interconnected systems can operate agnostically of whatever prime provided the system. These can vary in form factors and key design choices, but systems meet the same minimum thresholds.

A key question for program and portfolio officials is to what extent will they allow system competition to substitute for open hardware interfaces (component competition) and commonality (promoting reuse). A vendor may posit that vertical integration and rapid testing is their path to producibility. Alternately, vendors may push for designs with open interfaces that vary, which does not allow for easy reuse of items.

Acquire Mobilization Options

Under the MOSA "gray box" model, the government understands a module's function, behavior, and interfaces but does not hold the TDP that would allow for organic or contract-producer manufacturing.[48] An option to mobilize goes further, by including a franchising or manufacturing licensing agreement option clause or by arranging for a government reference design that can be directly distributed. This is a challenge to traditional DIB business models and thus may require considerable leverage to arrange.

The option to mobilize has two core advantages: addressing a wider base of manufacturing capacity and employing proven designs to minimize the difficulties of qualifying a new source. This approach draws on multiple concepts: the conversion of factories to wartime production during the Second World War, contract-production common in multinational supply chains, and new concepts such as manufacturing-as-a-service. However, the option to mobilize is more challenging than the disruptive but developed logic of MOSA gray boxes. The Department has rights to IP it fully funds, and MOSA benefits from government-purpose rights (GPR) to open interfaces. However, most development has mixed funding, and detailed manufacturing and process data is well protected absent an invocation of the Defense Production Act. Vendors seeking to ramp manufacturing often want to capture the process

[48] Taking the example of a lightbulb, the bulb must product light when 120VAC is applied and the interface is the familiar Edison 26 base. The guide notes "the consumer does not need to know the details of the manufacturing process used to extrude a Tungsten filament—nor does it care" (Guertin, 2025, p. 10).

knowledge that comes with doing it themselves. Contract-production is widely used in the commercial sector, but even contract-producers legendary for their protection of partner's IP, like the Tawain Semiconductor Manufacturing Corporation (TSMC), have seen partners become trillion-dollar companies, like NVDIA, or suffer setbacks in their core competencies, like Intel.

In the United States, scaling defense production for a given item typically involves mergers and acquisitions, spin-offs, joint ventures, a vendor expanding facilities, and/or international companies establishing U.S. subsidiaries. Contract-producers and facilities offering manufacturing as a service are standing up, but DIB use of this business model lags commercial use. Interestingly, the DIB does make more use of this model internationally despite higher regulatory barriers. Within the United States, contract-production would require a manufacturing license agreement from the IP holder, a TDP, and appropriate facility permits at the destination factory. For international production, tools, design, and manufacturing know how transfers will be regulated by the munitions list and commerce control list and will need appropriate licenses. These international arrangements often are part of the DIB deals to win foreign market access or U.S. government efforts to achieve foreign policy goals.

Due to the uncertain models, experimentation in this area is necessary and worthy of close study. One affordable mass model is franchise rights, under which the manufacturer agrees to arrange for new production facilities if surge capacity becomes needed. A variant on this approach would rely on special license arrangement options and the concept of IP as a service to prearrange the ability to surge production to a third party. Two ongoing efforts focus on bottleneck and difficult-to-qualify energetics subsystems. At AvMC, this effort included direct manufacturing at organic depots in parallel with industry. This approach pairs well with targeted commoditization to build up surge capability through more producible designs. Producible designs open the aperture for domestic commercial IB, the domestic organic IB, or in the forward IB of overseas depots and allied and partner production. In these examples, Pathfinders use a mix of direct funding of development and/or competitive pressure to gain rights and options that industry has hesitated to allow under other circumstances.

The speed with which CCA producers established production sites in Europe hints at how existing international co-production models could evolve to more rapidly boost production capacity and arrange for surge capability (Decker, 2025). The franchise model resembles the way producers establish or acquire subsidiaries across international borders. An option to mobilize could involve moving beyond company, program, or variant-specific facilities. Contract production and manufacturing as a service could have a large role in international production. One advantage of the United States in negotiating with international industry is that market access and the imprimatur of being qualified in the U.S. system can enhance the allied industries' ability to compete for international demand. International vendors may be willing to accept greater reliance on third-party production if qualification was part of the deal. Relatedly, there are more international vendors interested in production in the United States than can likely be maintained during a downturn, and each of these international vendors establishing their own facilities makes it harder to plan for economies of scale.

Conclusions

Pathfinding programs in the affordable mass space and beyond are experimenting with employing MOSA to achieve benefits for production capacity. MOSA guidebook recommendations on obsolescence are already relevant to building production agility to address challenges of boosting production capacity. Surge capability is an even greater challenge because demand is often correlated across programs and often across countries. Given the time and capital required to stand up new production lines, this means requiring (1) the U.S. government paying for or incentivizing the preservation of mothballed capacity in government or industry during downturns, (2) dual-use items or sufficiently producible military-spec items with necessary TDPs where military production can displace commercial production, (3) allied and partner production with more capacity than can be met by their domestic needs and that is flexible enough to shift to address bottlenecks and support surge. Pathfinders are most relevant to the second path but have relevant lessons for all three.

This paper offers an initial framework, summarized in

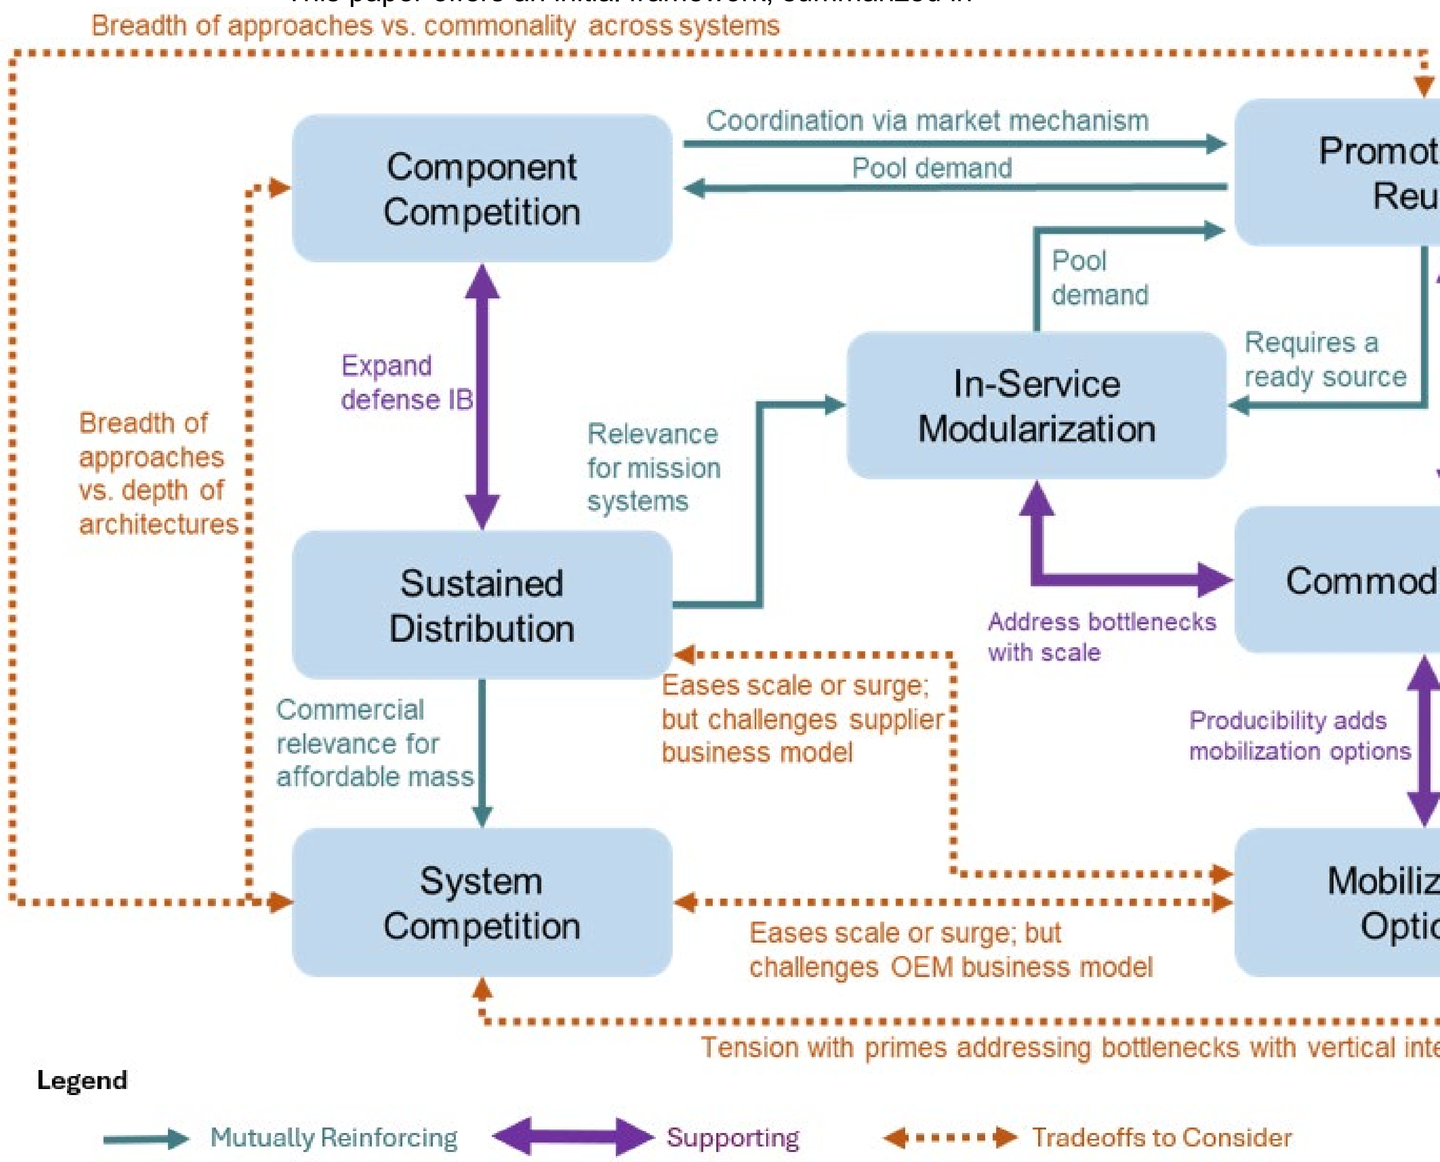

Figure 33, to assist in the consideration of tools that could accelerate production, especially with adoption at scale. The Joint Production Accelerator Cell (JPAC) will elevate and echo lessons learned from the challenges faced and successes learned from these efforts. The paper offers the following courses of action as ways of scaling lessons from pathfinder programs to accelerate munition production more broadly.

- Existing guidance on obsolescence on developing a product plan is also helpful for programs seeking to insure against production constraints writ large.

- Existing efforts on energetics subsystems suggest that MOSA can be part of the solution to accelerating the pace of qualification by lowering risk by commoditizing modules and acquiring mobilization options.
- Pathfinders in adjacent systems suggest that MOSA can aid in incorporating COTS and near-COTS to sustain markets that expand the munition IB.
- The Department has important leverage when ramping munition production. Both modularizing in-service munitions and acquiring mobilization options have production benefits, considered as part of negotiating a steady demand signal.
- Testing conformance of intersystem MOSA will aid in achieving and sustaining the benefits of system competition and ensuring the viability of on-ramps.

JPAC will build on this effort through research seeking to understand:

- What makes the affordable mass systems easier and cheaper to produce, besides them being only 80% to 90% effective compared to exquisite systems?
- What lessons can we learn that might be adaptable for those higher-end systems?
- How do we define standards and expectations for families of weapon systems that can enable multiple production lines or production done by different parties? What data rights, processes, and management properties are required to capture all of it in a TDP?

The experimentation by pathfinding systems affirms that MOSA is a combat capability relevant to addressing production constraints that limit the options available to operators even when budget is available. Consideration of MOSA, system competition, and the option to mobilize can play a supporting role in scaling production acceleration today and a vital role for future production agility, surge capability, and sustaining the munitions IB in the face of future budget downturns.

Appendix: Guided Conversation Questions

This ongoing project draws data from six guided conversations with munition and adjacent programs, subject matter experts, supplemented by desk and market research to identify core competencies of industry with assistance from relevant consortiums. The guided conversations asked experts and practitioners six questions:

1) ***Producibility***
    a. What, if any, overarching producibility concerns did you explicitly consider as part of your program or initiative?
    b. Are there any parts, components, or subsystems that are a heightened producibility concern for your program or initiative?

2) ***Architecture:***
    a. What, if any, key commercial standards, open standards, or government reference architectures are you employing or creating in your program or initiative?
    b. Were any of your architectural design or standard adoption decisions made with producibility as a critical factor?

3) ***Outcomes:***

a. Have you seen any producibility benefits of your architecture choices at your present stage in the life cycle of the program or initiative? Have any expected benefits for this stage not materialized?
b. Do you see any limits or risks related to affordable mass system adoption of a modular open system approach or government reference architectures?

Works Cited

Carroll, J. T. (2008, February). The reformers. *Air & Space Forces Magazine*.

Cook, C. (2023, March 14). *Reviving the arsenal of democracy: Steps for surging defense industrial capacity.* Center for Strategic and International Studies.

Decker, A. (2025, Juy 17). *General atomics plans robot wingman production for Europe*. Defense One. https://www.defenseone.com/business/2025/07/general-atomics-plans-robot-wingman-production-europe/406815/

Del Toro, C., Wormuth, C., & Kendal, F. (2024, December). *Modular open systems approach for Department of Defense weapon systems.* DoD. https://www.cto.mil/wp-content/uploads/2024/12/Tri-Service-Memo-Signed-17Dec2024.pdf

Fookes, R. (2023). *Air Force Materiel Command (AFMC) guidebook for implementing modular open systems approaches in weapon systems, Version 2.0.* Air Force Material Command.

Guertin, N. (2025). *Naval modular open system approach guidebook, Version 1.0.* DoD.

Gunzinger, M. A. (2021, November 2021). *Affordable mass: The need for a cost-effective PGM mix for great power conflict.* Mitchell Institute.

Hegseth, P. (2026). *2026 national defense strategy.* DoW.

Jones, S., & Palmer, A. (2024, March). *Rebuilding the arsenal of democracy.* Center for Strategic and International Studies.

Kendall, F. (2013). *Employment of Open Systems Architecture Contract Guidebook for Program Managers, Version 1.1.* Washington, DC: The Under Secretary of Defense for Acquisition, Technology, and Logistics.

Ludwigson, J. (2024, September 9). *GPS modernization: Delays continue in delivering more secure capability for the warfighter.* Government Affairs Office. https://www.gao.gov/products/gao-24-106841

OEDSE&A. (2024). *DoD producability and manufacturability engineering guide.* Office of the Under Secretary of Defense for Research and Engineering. https://www.cto.mil/wp-content/uploads/2024/06/DoD-Producibility-and-Manufacturability-2024.pdf

OSE&A. (2025). *Implementing a modular open systems approach in Department of Defense programs.* Office of the Under Secretary of Defense for Research and Engineering.

Pettyjohn, S. E. (2025, January 20). *Build a high–low mix to enhance America's warfighting edge and deter China.* Center for a New American Security.

Raytheon Company. (2011, June 20). *Raytheon Partners with NAMMO for second source of AMRAAM motors.* https://raytheon.mediaroom.com/index.php?item=1847

Sanders, G., & Aldisert, A. (2024, December). *Burden sharing via modular open system approaches: A collaborative path to affordable mass.* Center for Strategic and International Studies.

Shalom, M. (2026, January 26). *The primes aren't the real bottleneck in U.S. weapons production.* War on the Rocks. https://warontherocks.com/2026/01/the-primes-arent-the-real-bottleneck-in-u-s-weapons-production/

Wang, D. (2025). *Breakneck: China's quest to engineer the future.* W. W. Norton & Company.

Wood, P., & Stone, A. (2021). *China's ballistic missile industry.* China Arospace Studies Institute.

Acquisition Research Program
Department of Acquisition, Finance, and Manpower
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net