SYM-AM-26-136

Excerpt from the Proceedings

of the

Twenty-Third Annual Acquisition Research Symposium and Innovation Summit

Volume III

"Accelerating Warfighting Capabilities"

The AI Acquisition Nexus: A Framework for Program Managers in the U.S. Department of War

Published: April 30, 2026

The research presented in this report was supported by the Acquisition Research Program, Graduate School of Defense Management at the Naval Postgraduate School.

To request defense acquisition research, please contact:

Acquisition Research Program
Department of Defense Management
Naval Postgraduate School
E: arp@nps.edu
www.acquisitionresearch.net

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

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

The AI Acquisition Nexus: A Framework for Program Managers in the U.S. Department of War

LT Sean Courtney, USN—Engineering Duty Officer. [sean.t.courtney.mil@us.navy.mil]

Abstract

The Department of War (DoW) has issued an unambiguous mandate for wartime speeds of AI adoption. Secretary Hegseth's January 9, 2026, memorandum directs the department to become an "AI-first warfighting force" through seven Pace-Setting Projects (PSPs), aggressive data access, decrees, and a 30-day AI Model Parity requirement. Simultaneously, the FY2026 National Defense Authorization Act (NDAA) establishes binding legislative requirements for AI cybersecurity frameworks, cross-functional assessment teams, and AI sandbox environments. These directives are consequential, yet both the memorandum and the NDAA are primarily top-down instruments. They establish timelines and deliverables but do not resolve the ground-level acquisition gaps that have persisted across service branches since the Government Accountability Office (GAO) first documented them in 2023. This report conducts a comprehensive analysis of the AI acquisition methods employed by the U.S. Army, Navy, and Air Force, mapping their strengths and structural limitations against the requirements of the new interoperability, Test and Evaluation (T&E) standardization, intellectual property management, workforce readiness, and AI cybersecurity. Where the 2026 directives already address an identified gap, this report proposes specific implementation mechanisms to operationalize those mandates at the service level. Where genuine gaps remain that neither the memorandum nor the NDAA addresses, this report offers nine recommendations, including a tiered assurance certification pathway, an AI Bill of Materials (AI-BOM) standard, and the integration of adversarial security testing into the operational T&E lifecycle. The framework proposed herein is designed not to replace service-level approaches but to provide the mending needed to transform isolated experiments into joint warfighting capabilities.

Introduction

The reorganization of the Department of Defense into the Department of War in late 2025 was a declarative signal of national intent designed to realign the institutional mindset of the American military establishment from sustained deterrence to active competitive engagement with near-peer adversaries. Within this context, the integration of Artificial Intelligence has been elevated from a modernization priority to a wartime necessity.

Secretary Hegseth's January 9, 2026, memorandum codified this urgency. The memo declares that "AI-enabled warfare and AI-enabled capability development will re-define the character of military affairs over the next decade" and that "this transformation is a race" (Secretary of War, 2026, p. 1). It directs the department to "accelerate America's Military AI Dominance by becoming an 'AI-first' warfighting force across all components, from front to back" (Secretary of War, 2026, p.1). Yet, this vision must be executed by an acquisition workforce operating within a framework that, as the GAO has documented, lacks "department wide guidance for how its components should approach acquiring AI" (U.S. Government Accountability Office [GAO], 2023, p. 1).

The resulting actions taken are paradoxical. The strategic leadership demands wartime speed, while the acquisition workforce operates at a peacetime tempo. The January 2026 memorandum and the FY2026 NDAA collectively establish the most consequential AI directives in the department's history, but neither resolves the ground-level execution gaps that persist across service branches. This report seeks to bridge that distance by analyzing current branch-specific methods, mapping them against the new mandates, identifying where those mandates already address known gaps, and proposing any solutions where they do not.

The Policy Landscape

Brief History

The DoW's AI posture is built on years of policy evolution that began with the U.S. response to China's 2017 New Generation AI Development Plan. The National Security Commission on AI's 2021 final report warned that "the U.S. was not prepared to defend or compete in the AI era" and called for federal AI R&D investment to double annually, reaching $32 billion by 2026 (National Security Commission on Artificial Intelligence [NSCAI], 2021). The DoD's 2020 adoption of five Ethical AI Principles, which was grounded in "the law of war, the U.S. Constitution, and Title 10 of the U.S. Code" (DoD, 2020b, para. 3), and the 2022 Responsible AI (RAI) Strategy signed by Deputy Secretary Hicks established the ethical and governance foundation. The Chief Digital and Artificial Intelligence Officer (CDAO), codified in DoD Directive 5105.89 as the "principal officer responsible for accelerating the adoption of data, analytics, and AI capabilities … across the defense enterprise" (DoD, 2024, sec. 1), was designated as the central governance authority.

That previous work defines RAI as a "force multiplier" and explicitly rejects the notion that it can be achieved through a "simple 'ethical review' prior to acquisition or a final check-in-the-box before fielding" (Hicks, 2022, p. 4). However, such a framework was designed for a department in peacetime. The 2026 mandates demand trust at speed, a fundamentally different operational requirement.

Call for AI Dominance (January 2026 Memorandum)

The Secretary of War's memorandum represents a complete shift from principles to timelines. Its central directive establishes seven PSPs spanning warfighting, intelligence, and enterprise mission areas. They include the following:

- *Swarm Forge* - competitive discovery and scaling of AI-enabled warfighting concepts
- *Agent Network* - AI-agent development for battle management and decision support
- *Ender's Foundry* - AI-enabled simulation capabilities
- *Project Grant* - dynamic deterrence transformation
- *Open Arsenal* - accelerating technical intelligence (TechINT)-to-capability pipelines
- *GenAI.mil* - democratized AI access at all classification levels
- *Enterprise Agents* - AI agent deployment for enterprise workflows

The PSPs are framed as "tangible, outcome-oriented vehicles for rapidly completing our buildout of the foundational AI enablers needed to accelerate AI integration across the entire Department" (Secretary of War, 2026, p. 1).

The memorandum establishes several binding operational directives beyond PSPs. On data access, it directs all military departments and components to "deliver their current catalogs" to the CDAO "within 30 days of the date of this memorandum" and authorizes the CDAO to "direct release of any DoW data to cleared users with valid purpose" (Secretary of War, 2026, p. 3). On authorization velocity, the Under Secretary of War for Research and Engineering (USW[R&E]) is directed to "establish a monthly 'Barrier Removal Board' with authority to waive non-statutory requirements and escalate blockers for immediate resolution" (Secretary of War, 2026, p. 4). On model currency, the CDAO is directed to "establish a delivery and integration cadence with AI vendors that enables the latest models to be deployed within 30 days of public release" as "a primary procurement criterion" (Secretary of War, 2026, p. 4). On architecture, program managers are directed to "enforce Modular Open System Architectures (MOSA) along

with the ‘DoD Data Decrees,’ exposing modular interfaces and associated documentation sufficient for third-party integration without prime contractor support” (Secretary of War, 2026, p. 5).

The memo also considers cultural impact on a broad scale. It declares that “the risks of not moving fast enough outweigh the risks of imperfect alignment” and mandates that “exercises and experiments that do not meaningfully incorporate AI and autonomous capabilities will be reviewed by the Director of Cost Assessment and Program Evaluation for resourcing adjustment” (Secretary of War, 2026, pp. 4–5). The CDAO is redesignated as a “Wartime CDAO” with expanded authorities for cross-domain data access and rapid authority to operate (ATO) reciprocity (Secretary of War, 2026, p. 4).

Foundations for AI Security and Governance (The FY2026 NDAA)

The FY2026 National Defense Authorization Act, signed December 18, 2025, provides the legislative complement to the executive memorandum. While the secretary’s memo emphasizes speed, the NDAA imposes requirements for security, governance, and accountability that the department must satisfy.

Section 1512 directs the DoD to “develop, within 180 days, a department-wide policy for cybersecurity and governance of AI and machine learning systems, addressing AI/ML-specific threats, lifecycle cybersecurity measures, industry frameworks, governance standards, and workforce training” (NDAA, 2026, §1512). The threats enumerated include model serialization attacks, model tampering, data leakage, adversarial prompt injection, model extraction, model jailbreaks, and supply chain attacks. Section 1513 requires a comprehensive risk-based framework for cybersecurity and physical security standards for AI/ML technologies, to be implemented “as an extension or augmentation of existing cybersecurity frameworks developed by the Department of Defense,” including the Cybersecurity Maturity Model Certification (CMMC) program (NDAA, 2026, §1513). Section 1533 tasks the secretary with establishing a cross-functional team for AI model assessment and oversight by June 2026, which must develop a department-wide assessment framework including “standards for performance of AI models, testing procedures, security requirements and principles for the ethical use of AI” (NDAA, 2026, §1533). Additional provisions mandate AI sandbox environments (§1534), an AI Futures Steering Committee (§1535), and revised mandatory cybersecurity training to include AI-specific content (§1515).

Taken together, the January memorandum and the NDAA establish the most comprehensive set of AI directives in the department’s history. The memo provides the speed and direction; the NDAA provides the security and governance guardrails. The challenge for the acquisition workforce is that these two must be executed simultaneously, and neither provides the service-level implementation guidance necessary to do so.Current AI Acquisition Methods Across the Services

The U.S. Navy

The Navy is perhaps the most mature AI adopter of all the service branches, pursuing a strategy of foundational architecture and rapid experimentation. The Program Executive Office (PEO) IWS X program’s mission to “converge to a common Integrated Combat System (Hardware/Software) across all surface combatants” (PEO IWS, 2022, p. 1) provides the architectural backbone through its government-owned software factory (“The Forge”) and hardware counterpart (“The Foundry”), delivering Infrastructure-as-a-Service and Platform-as-a-Service for future AI deployment across ship classes.

Task Force 59, established in September 2021 under 5th Fleet, represents the Navy’s operational experimentation. TF 59 deployed commercially sourced unmanned surface vessels

with AI-driven anomaly detection to build a “digital ocean” of persistent surveillance (Vincent, 2023). The task force deliberately circumvented traditional acquisition norms, integrating commercial technologies through exercises like Digital Horizon and dramatically compressing the feedback cycle between operational need and technological solution. Naval Information Warfare Systems Command (NAVWAR)’s Project Overmatch completes the Navy’s trifecta, building Joint All-Domain Command and Control (JADC2) network architecture and employing new methods like the AINetANTX Prize Challenge, which gave participants access to government datasets through the Overmatch Software Armory (Gamboa, 2021).

The Navy’s primary limitation is that its ecosystem is heavily Navy-centric. The Forge, Project Overmatch, and TF 59’s operational concepts were designed for maritime warfare. There is no inherent mechanism for translating these capabilities across service boundaries without significant re-engineering.

The U.S. Air Force

The Air Force has distinguished itself through formalized acquisition methodologies. Its AI Acquisition Guidebook, developed with MIT, provides the most structured process available to any service PM, walking users from problem definition through contract strategy and deliverable specification (Department of the Air Force & MIT, 2022). The guidebook leverages the MITRE Innovation Tool Kit’s Problem Framing Canvas to help teams “become intrinsically aware of the various base factors causing the capability gap” before seeking solutions (Department of the Air Force & MIT, 2022, p. 12). Operationally, the Air Force has pursued AI across predictive maintenance (Condition-Based Maintenance Plus [CBM+]), intelligence analysis, and autonomous systems through programs like Skyborg and the Collaborative Combat Aircraft.

The Air Force’s weakness is that its exemplary methodology remains an informal guidebook rather than an instruction. Individual program offices adopt it inconsistently for this reason. More critically, the Air Force’s ATO process for certifying software on its networks routinely requires three to six months, making it directly incompatible with the 30-day Model Parity mandate.

The U.S. Army

The Army’s AI acquisition approach is shaped by its network modernization strategy and its organizational structure, which is oriented around platform-specific Program Executive Offices (PEO Soldier, PEO Ground Combat Systems, PEO Aviation, and others). This structure means that AI capabilities entering the Army acquisition pipeline tend to do so as components of existing platform programs rather than as standalone capabilities with dedicated acquisition strategies. When an AI-driven targeting aid is developed for a ground combat vehicle, for example, it is typically managed within the PEO responsible for that vehicle, with the acquisition strategy, contracting apparatus, and T&E plan shaped by the platform’s program rather than by AI-specific considerations. This creates a structural risk that AI capabilities are developed within the stovepipes of individual PEOs, leading to potential redundancy and incompatibility between implementations that serve similar functions across different platforms.

The Army has made significant efforts to address this fragmentation. Task Force Lima, established to evaluate generative AI applications for military use, has provided a rapid experimentation method that cuts across those PEO boundaries. Project Linchpin, managed by Army Futures Command, aims to provide a centralized cloud-based AI/ML development and deployment infrastructure intended to standardize tools and processes for developing AI models across Army programs. Both efforts represent important steps toward a more coherent AI acquisition posture. However, Project Linchpin faces integration challenges with the Army’s existing network infrastructure. The Army’s ongoing transition from the legacy Warfighter

Information Network-Tactical (WIN-T) to the Integrated Tactical Network (ITN) under the Network Cross-Functional Team means that network modernization and AI integration must proceed in parallel.

The Army's ATO process is also noteworthy. While all three services face authorization delays, the Army's process is complicated by the diversity of its computing environments, ranging from enterprise cloud systems at garrison to disconnected tactical edge devices in the field. An AI capability that achieves ATO on the Army's enterprise network does not automatically carry authorization to the field, requiring a separate assessment for each deployment context. This adds layers of certification that the Navy and the Air Force do not face to the same degree.

Assessment

The Defense Acquisition University (DAU; 2025) *Program Management Fundamentals Handbook* acknowledges that the AAF "requires a great deal of creative and critical thinking by the project/program managers and their teams" and is not just a checklist (p. 6). This flexibility becomes a liability when AI demands standardization across service boundaries. Each service has developed genuine innovation, but those innovations are not integrated. The Navy's Development, Security, and Operations (DevSecOps) pipelines, the Air Force's problem-framing toolkit, and the Army's platform-integration experience represent strength that has never been harmonized into a joint capability - and it needs to.

Gaps Between the Service-Level Methods and the 2026 Mandate

Authorization Time

The Model Parity mandate presumes an authorization process that can accept, validate, and clear an AI model for deployment on a new service's network within 30 days. Current ATO processes across all three services routinely require three to six months or longer. The January memorandum partially addresses this through the Barrier Removal Board, which can "waive non-statutory requirements and escalate blockers" (Secretary of War, 2026, p. 4), and through the Wartime CDAO's authority for "rapid ATO reciprocity" (Secretary of War, 2026, p. 4). However, the Barrier Removal Board is a reactive mechanism that addresses blockers case-by-case after they arise. It does not establish a proactive, standing certification pathway that enables cross-service model acceptance by default.

Data Interoperability

The January memo's data access directives are among its strongest provisions, requiring federated data catalogues within 30 days and authorizing the CDAO to direct release of any DoW data. These directives solve the data discoverability problem. And for the case of deploying a model as-is from one service to another, the model itself is portable and the input pipeline can be documented in a Model Card.

The point of friction occurs in two specific scenarios. First case, when a receiving service needs to retrain or fine-tune a model for its own operational environment. A Navy model trained on maritime sonar data is not useful to the Army without Army-relevant training data, and if the Army's datasets are labeled with different conventions, stored in different formats, or governed by different access controls, the fine-tuning pipeline breaks at the data-engineering level. Second case, when the T&E community at the receiving service needs to independently validate the model. DoDM 5000.101 requires that government test datasets be "independent of the contractor's training data to provide a true, unbiased assessment of the model's generalization capabilities" (Director, Operational Test & Evaluation, 2024, sec. 3). If the receiving service's T&E team must validate a transferred model using their own test data, and that data follows different labeling conventions, it must be reformatted before testing can proceed. The Data Card concept in DoDM 5000.101 provides an excellent documentation

framework, but without binding standards for how Data Cards are structured and what labeling conventions they enforce, they describe datasets that remain functionally incompatible across services for retraining and validation purposes. Standardization of Data Cards and Model Cards is therefore the highest-leverage intervention in this space.

T&E Standardization and the Assurance Case Divergence

DoDM 5000.101's explicit goal is to "improve test planning, test rigor, and implementation of leading practices for identifying and quantifying [the] risks of AI-enabled systems" (Director, Operational Test & Evaluation, 2024, p. 1). The emerging "assurance case" concept, defined as "a structured argument, [supported by evidence from T&E], that a system is sufficiently trustworthy [to be fielded] in a specific range of operational contexts" (Developmental Test, Evaluation, and Assessment [DTE&A], n.d., para. 2), represents the gold standard for AI accreditation.

It is appropriate and expected that different operational applications require different assurance cases. A Navy AI for submarine sonar classification and an Air Force AI for aerial tracking serve different operational contexts and should be evaluated accordingly. The divergence this report identifies is not about application-specificity, which is proper, but about evidentiary standards. Each service's T&E authority currently sets its own thresholds for what constitutes sufficient evidence in an assurance case, including the statistical methodologies considered valid, the sample sizes required for robustness demonstrations, and the adversarial test scenarios deemed necessary at comparable risk levels. There is no joint standard governing the format of assurance arguments, the thresholds for acceptable risk, or the test methodologies recognized as valid across services. This means that even for functionally equivalent applications, a model certified by one service's T&E authority carries no presumptive weight with another's.

The practical consequence is significant. Consider a predictive maintenance AI certified by the Air Force for aircraft engine monitoring. If the Navy wished to adapt it for shipboard turbine monitoring, a functionally similar application at different scale, the Navy's T&E authority would conduct an entirely independent evaluation with no obligation to accept or even review the Air Force's assurance case. The Air Force's test data, methodology, and conclusions carry no formal weight. The model must effectively prove itself from scratch, even when the operational context is closely comparable. DoDM 5000.101 organizes testing into AI Model T&E, Systems Integration T&E, and Operational T&E, but each service implements these stages through its own institutional process. The NDAA's Section 1533 requirement for a cross-functional assessment team may eventually address this, but its output of a framework by June 2027 arrives too late for the PSP timelines already underway.

Intellectual Property Across Service Boundaries

The GAO (2023) concluded that the DoD, by failing to issue AI-specific acquisition guidance, "is missing an opportunity to ensure that it is consistently acquiring AI capabilities in a manner that accounts for the unique challenges associated with AI" (p. 18). The January memo's MOSA directive and instruction to incorporate "any lawful use" language into future AI contracts (Secretary of War, 2026, p. 5) addresses the vendor-to-government IP relationship for prospective acquisitions.

In principle, Government Purpose Rights (GPR) granted under a DoW contract should permit use across all components, as the service falls under a single governmental authority. In practice, the friction arises in two areas once again. First case, many AI contracts involve commercial items acquired under Federal Acquisition Regulation (FAR) Part 12, where the government often accepts commercial license terms rather than negotiating GPR. Commercial licenses can restrict use to specific programs, users, or environments. Second case, contracts

predating the January memo may define IP deliverables in reference to a specific service's program of record, creating ambiguity about whether use in a different service's operational environment exceeds the negotiated scope. The memo's "any lawful use" and MOSA directives are designed to resolve this going forward, but they do not retroactively cure the existing contract base. For legacy contracts, cross-service model transfer may require contract-by-contract legal review unless a prospective Joint Deployment Rights standard is established.

Workforce Readiness

The memo directs each component to provide "AI hiring and talent development plans to the Under Secretary of War for Personnel and Readiness within 60 days" (Secretary of War, 2026, p. 3). The NDAA's Section 1515 required revised cybersecurity training to include AI content within one year. These provisions address the talent pipeline but not the cross-service fluency problem. Each service's acquisition workforce brings its own institutional vocabulary, risk tolerance, and process assumptions. The RAI Toolkit and DAU training resources provide a starting point but have not been adopted uniformly. Without a mandated joint training curriculum that establishes a common operational vocabulary for AI acquisition, the workforce gap will continue to impede the standardization that the Model Parity mandate and the PSPs require.

AI Cybersecurity as an Integrated Discipline

The FY2026 NDAA's Sections 1512 and 1513 mandate both a department-wide cybersecurity policy for AI and a comprehensive risk-based security framework. The threats enumerated are specific and technically serious, including data poisoning, model tampering, adversarial prompt injection, model jailbreaks, model extraction, and supply chain attacks (NDAA, 2026, §1512). Section 1513 specifies that the framework must address "workforce risks, including insider threat risks, training and workforce development requirements regarding artificial intelligence security awareness, [AI]-specific threats and vulnerabilities," as well as tampering risks and unintended data exposure (NDAA, 2026, §1513). The NDAA further directs that this framework be implemented "as an extension or augmentation of existing cybersecurity frameworks," including CMMC, signaling Congress' intent that AI security requirements will eventually flow down through the defense industrial base in the same manner as existing cybersecurity certifications (NDAA, 2026, §1513).

The National Institute of Standards and Technology's (NIST's) recently updated AI 100-2 (Adversarial Machine Learning) provides the most comprehensive taxonomy of AI-specific attack vectors available, covering both predictive and generative AI systems. The taxonomy organizes threats by lifecycle stage and attacker objective. Attacks done at training time include data poisoning, where adversaries contaminate training datasets to introduce hidden vulnerabilities or systematic misclassifications. Deployment-time attacks include two vectors: (1) evasion attacks, where carefully crafted inputs cause a trained model to produce incorrect outputs, and (2) privacy attacks, such as membership inference and data reconstruction, where an adversary extracts sensitive information about the training data by interacting with the deployed model. For generative AI specifically, prompt injection attacks represent a distinct threat class in which adversaries manipulate system behavior through malicious inputs. The NIST guidance acknowledges that "at this stage with the existing technology paradigms, the number and power of attacks are greater than the available mitigation techniques" (NIST, 2025, p. x), an honest assessment that underscores the immaturity of AI security as a discipline and the urgency of integrating it into the T&E process.

However, neither the NDAA nor the January memo integrates adversarial security testing into the operational T&E lifecycle established by DoDM 5000.101. The NDAA's security framework (due within 180 days of enactment) and DoDM 5000.101's T&E process exist as parallel tracks. An AI system could pass the manual's assurance case for operational

performance, including robustness testing against out-of-distribution inputs, while remaining vulnerable to deliberate adversarial attacks that exploit specific weaknesses in the model architecture. This parallel-track problem is the most urgent integration challenge facing the acquisition workforce, because it means that the department could field AI systems that are operationally competent yet adversarially fragile.

Implementation Methods and Recommendations

Recommended Implementations - Mechanisms for Gaps

Issue a Joint AI Acquisition Instruction (DoWI 5000.XX)

The GAO (2023) recommended that the secretary "prioritize [the establishment of] department-wide AI acquisition guidance" (p. 26). The January memorandum provides strategic direction but not a procedural framework. A binding DoW Instruction should codify the memo's directives into acquisition requirements, including mandatory Data Cards and Model Cards as contract deliverables, MOSA compliance verification checkpoints, the 30-day data catalogue compliance mechanism, and required use of the Problem Framing Toolkit before any AI solicitation. This translates the memo's directives into the procedural language that contracting officers and PMs require for consistent execution. It could be accomplished within 90 days.

Operationalize Data Access Through a Federated Data Commons

The memo's catalogue requirement solves discoverability but not interoperability for retraining and validation. A Federated Joint Data Commons (FJDC) built on the Overmatch Software Armory model should provide a governed architecture where each service maintains data sovereignty while making datasets accessible through a common API. Critically, the FJDC must enforce binding Data Card standards with common labeling conventions and metadata schemas so that datasets transferred for retraining or independent T&E validation are compatible at the data-engineering level without reformatting. Initial operating capability should be achievable within 180 days by leveraging cloud investments.

Establish a Joint AI Acquisition Workforce Certification

The memo's 60-day talent plan requirement and the NDAA's Section 1515 training mandate addresses supply and awareness, but not cross-service fluency. A mandatory Joint AI Acquisition Certification, administered through DAU, should establish the common vocabulary and process understanding that the PSPs require. The curriculum should cover AI technical fundamentals, the DoW ethical principles and RAI Toolkit, the new acquisition instruction, Data Card and Model Card development, adversarial AI security concepts from NIST AI 100-2, and cross-service interoperability standards. Inception achievable within 90 days and required for all personnel in AI acquisition billets before year's end.

Recommended Actions - Addressing Gaps Remaining

Establish a Tiered Joint AI Certification Pathway

The Barrier Removal Board is highly reactive. What is needed is a proactive, standing mechanism that calibrates certification rigor to risk. This report proposes a three-tier Joint AI Model Certification Pathway administered by a Joint AI Model Certification Board (JAMCB), chaired by the Wartime CDAO. Tier 1, for low-risk enterprise AI (chatbots, document summarization, administrative automation), would follow a 72-hour automated certification process based on standardized security scans and performance benchmarks. Tier 2, for mission-support AI (intelligence analysis aids, logistics optimization, predictive maintenance), would follow a 15-day reciprocal review process where a model certified by one service undergoes a streamlined integration assessment by others. Tier 3, for safety-critical AI (autonomous weapons, targeting, combat system integration), would follow the full DoDM

5000.101 assurance case pathway with joint authority, rather than service-specific accreditation. This tiered approach reconciles the memo's demand for speed with DoDM 5000.101's demand for rigor by matching the process to the risk.

Integrate Adversarial Security Testing Into the T&E Lifecycle

Neither the NDAA's security framework mandates nor DoDM 5000.101 currently require adversarial red-teaming as a formal phase in AI Model T&E. This report recommends that adversarial security testing, informed by the NIST AI 100-2 taxonomy and the NDAA Section 1512 threat enumeration, be incorporated as a mandatory gate in the AI Model T&E stage of every AI acquisition program. This testing should evaluate resilience against data poisoning, evasion attacks, model extraction, prompt injection for generative systems, and supply chain integrity of pre-trained model components. Results should be documented in a standardized section of the Model Card, creating a unified artifact that captures both operational performance and security posture. This eliminates the parallel-track problem identified in Gap F by fusing security assurance into the same process that evaluates operational trustworthiness.

Standardize an AI Bill of Materials (AI-BOM)

The NDAA's Joint Explanatory Statement indicates that software bill of materials (SBOM) policies "should also apply, where feasible, to AI systems used, developed, or acquired by DoD" (NDAA, 2026, Joint Explanatory Statement). This report proposes operationalizing that guidance through a standardized AI Bill of Materials that packages seven components into a single, auditable artifact.

The AI-BOM should contain the following elements:

1. The Model Card documenting architecture, provenance (whether built from scratch or adopted from a pre-trained model), performance metrics, known biases, and robustness characteristics.
2. The Data Card documenting training and test dataset provenance, coverage, labeling methodology, and known gaps.
3. The training provenance chain documenting the complete lineage of how the model was produced, including hyperparameters, training infrastructure, and any intermediate checkpoints.
4. A pre-trained component manifest identifying all foundation models, transfer-learned components, or third-party model weights incorporated into the system, with their respective licenses and origins.
    a. This element is particularly important because publicly available pre-trained models often carry license restrictions (such as CC-BY-NC 4.0 for certain popular checkpoints) that are incompatible with commercial or government deployment.
5. A software dependency manifest identifying all libraries, frameworks, and runtime dependencies.
6. The adversarial security test results documenting the outcomes of red-teaming conducted under the integrated T&E process proposed in Recommendation B2.
7. The IP rights allocation documenting the specific government rights (GPR, Unlimited Rights, or commercial license terms) that apply to each component.

The AI-BOM would travel with every AI model transferred between services under the Model Parity mandate, providing the receiving service with everything needed for rapid assessment under the tiered certification pathway:

- *Tier 1 Certifications* - the AI-BOM enables automated security scans and compliance checks.
- *Tier 2 Reciprocal Reviews* - provides the receiving service's T&E team with the evidentiary foundation for a streamlined integration assessment.
- *Tier 3 Full Assurance Cases* - provides the documentation backbone upon which the joint accreditation authority builds its structured argument.

The AI-BOM also makes cross-service deployment rights transparent at the point of transfer by explicitly documenting IP allocations at the model level, reducing the need for contract-by-contract legal review of legacy agreements.

Mandate Cross-Service IP Rights at Contract Formation

The memo's MOSA directive and "any lawful use" language addresses the vendor-to-government relationship for future contracts. For the existing portfolio of commercial AI contracts that may restrict cross-service use, a prospective solution is needed. This report recommends that the Wartime CDAO issue a directive requiring all future AI acquisition contracts to include a standard "Joint Deployment Rights" clause granting GPR for use, modification, testing, and redeployment across all DoW components. This clause should be developed by the Under Secretary of War for Acquisition and Sustainment in coordination with service General Counsels and incorporated into the Defense Federal Acquisition Regulation Supplement (DFARS) within 120 days. For existing contracts, the CDAO should direct a portfolio review to identify AI assets with cross-service restrictions and prioritize renegotiation for those assets most relevant to the PSPs.

Establish Post-Deployment Governance at Wartime Speed

DoDM 5000.101 requires that PMs have processes "in place to detect, track, and respond to deviations in performance" post-deployment (Director, Operational Test & Evaluation, 2024, sec. 4). The January memo's monthly reporting cadence focuses on new capabilities instead of fielded ones drifting out of specification. This report recommends that all AI programs of record establish automated performance monitoring pipelines with drift detection integrated into the CDAO's reporting dashboard. Quarterly joint T&E reviews for all PSP-associated capabilities and annual assurance case updates for all fielded AI systems should be mandatory. Operational feedback mechanisms modeled on TF 59's rapid experimentation cycle should route user insights directly to development teams. This ensures the fielded portfolio remains trustworthy as operational environments and adversary tactics evolve. The urgency of this recommendation is reinforced by the GAO's (2023) earlier finding that the DoD lacked guidance defining outcomes and monitoring accountability for AI-related activities, including AI-specific key performance indicators, a recommendation with which the department concurred but had not implemented as of 2023 (p. 10).

Designate PSP Technical Standards as Precedent

The seven PSPs will, by necessity, establish de facto joint standards for data formats, API interfaces, model packaging, AI-BOM structure, and T&E criteria. The Wartime CDAO should formalize this dynamic via directive, establishing that technical standards from any PSP are automatically elevated to DoW-wide standards unless a service formally objects within 30 days with a technically substantiated alternative. The PSPs that most directly force this standardization are:

- *Swarm Forge* - requires multi-domain interoperability across all six gaps simultaneously
- *Agent Network* - requires a common intelligence data ingestion layer and serves as the proving ground for standardized Data and Model Cards

- *GenAI.mil* - provides a common platform that can streamline the authorization problem by enabling platform-level ATO rather than model-level ATO

This "standards-by-doing" approach takes advantage of operational momentum to drive the combined effort faster than any traditional standards process could achieve.

Cost Considerations and Risk Assessment

Establishing a Joint AI Acquisition Instruction, building a Federated Data Commons, standing up a JAMCB, and developing a workforce certification program require dedicated funding, personnel, and organizational bandwidth. These represent a deliberate front-loaded investment. The FJDC alone will require cloud infrastructure procurement, API development, and the establishment of a data governance team within the CDAO's office. The JAMCB will require dedicated billets from each service's T&E community. The workforce certification program will require curriculum development, instructor faculty, and temporary reductions in workforce capability as personnel cycle through training. None of these costs are trivial. This report does not attempt to provide quantitative cost estimates for the proposed framework, and that absence is a limitation. Order-of-magnitude costing for the FJDC, the JAMCB, and the certification program should be a first-order task for the Wartime CDAO's implementation team, informed by analogous investments in existing joint infrastructure.

However, the alternative is not net-zero cost. Continuing with the fragmented status quo means continued investment in redundant AI development efforts across services, each building similar capabilities independently. It means vendor lock-in premiums on contracts that failed to negotiate adequate data and IP rights. It means the cost of programs ultimately rejected by operators because insufficient T&E failed to calibrate trust appropriately. And it means the strategic cost of fielding AI capabilities too slowly to provide operational advantage in a competitive environment where, as the secretary's memo states, "speed wins" (Secretary of War, 2026, p. 4). The framework proposed here is designed to prevent these downstream negative costs.

The expansion of the wartime CDAO's authority carries the risk of over-centralization. If the JAMCB becomes a slow-moving approval body, or if the FJDC's governance is overly restrictive, centralization becomes a bottleneck. Mitigation requires embedding agile principles into every new institution, with strict decision timelines, automatic approval provisions for Tier 1 certifications, and a culture that treats speed as a feature rather than a risk. The tiered certification pathway is specifically designed to prevent the JAMCB from applying Tier 3 rigor to Tier 1 problems.

A related but distinct risk is institutional resistance. The services have deep organizational incentives to retain independent T&E certification authority, and the Navy's T&E community in particular is unlikely to accept a JAMCB ruling that an Army-certified model meets maritime operational standards without significant forcing mechanisms. History offers few examples of successful cross-service certification bodies that did not face sustained bureaucratic resistance in their early years. To mitigate this, the JAMCB should include mandatory service representation with rotating vice-chairs, ensuring that no single service perceives the board as an external imposition. Additionally, a sunset clause that returns certification authority to individual services if the JAMCB fails to meet its own published decision timelines within the first 18 months would create accountability in both directions, incentivizing the board to deliver on its speed commitments while giving the services a credible assurance that centralization will not become permanent dysfunction.

The emphasis on government data rights, cross-service IP clauses, and transparency requirements will create friction with some commercial vendors. The 2022 National Defense

Strategy acknowledged this, stating that the DoD (2022) "will be a fast-follower where market forces are driving commercialization of military-relevant capabilities in trusted artificial intelligence and autonomy" (p. 9). The January memo reinforces this posture by mandating engagement with "America's world-leading companies" through "creative partnerships" (Secretary of War, 2026, p. 3). The Joint Deployment Rights clause proposed in Recommendation B4 will narrow the vendor pool for legacy-style proprietary contracts, but it filters for partners willing to support full-lifecycle collaboration. The Brennan Center for Justice has cautioned that acquisition reforms broadening commercial exemptions "may reduce transparency and weaken DoD oversight" (Toh & Gledhill, 2024, para. 3). The tiered certification pathway and the AI-BOM are designed to maintain appropriate oversight while enabling the speed the memo demands.

Conclusion

The DoW possesses, for the first time, both the strategic direction and the legislative authority to transform AI from a collection of service-level experiments into a decisive joint warfighting capability. The January 2026 memorandum establishes the velocity. The FY2026 NDAA establishes the guardrails. The seven PSPs provide the vehicles for execution. What remains missing is the implementation level between them.

The six gaps identified in this analysis are structural impediments to the stated objective of AI-enabled dominance. The authorization velocity gap means the 30-day mandate cannot be met by existing processes. The data standardization gap means federated catalogues enable discovery but not interoperability for retraining and validation. The T&E evidentiary divergence means certification by one service carries no weight with another service. The IP gap means cross-service model transfer faces legal ambiguity in the legacy contract portfolios. The workforce gap means services lack the common vocabulary to execute joint programs. And the cybersecurity gap means AI systems could be operationally certified yet adversarially vulnerable.

The nine recommendations address these gaps through complementary approaches. Three operationalize existing directives by providing the procedural mechanisms that the memo and NDAA mandate but do not specify. Six offer new insights, including the tiered certification pathway that calibrates rigor to risk, the integration of adversarial security testing into the T&E lifecycle, the AI Bill of Materials as a portable assurance artifact, cross-service IP rights at contract formation, post-deployment governance at wartime pace, and standards-by-doing through PSP precedent.

The memo is correct that "the risks of not moving fast enough outweigh the risks of imperfect alignment" (Secretary of War, 2026, p. 4). But speed without structure produces chaos, not advantage. The framework proposed here enables speed with sufficient structure, calibrating rigor to risk so that low-consequence AI moves at wartime speed while high-consequence AI receives the scrutiny that warfighter safety demands. The acquisition workforce has the strategic mandate it has long needed. What it needs now are the tools to execute, and there is little time to make mistakes.

References

Chief Data Officer. (2021). *DoD data stewardship guidebook.* U.S. DoD. https://www.ai.mil/Portals/137/Documents/Resources%20Page/DoD%20Data%20Stewardship%20Guidebook.pdf

Chief Digital and AI Officer. (n.d.). *Pathway to AI readiness: Data quality.* U.S. DoD. https://www.ai.mil/About/Resources/Pathway-to-AI-Readiness/Data-Quality/

Chief Digital and AI Officer. (2023). *Responsible AI (RAI) toolkit.* U.S. DoD. https://www.ai.mil/Initiatives/Responsible-AI/

Chief Digital and AI Officer. (2024). *Operational test and evaluation of artificial intelligence-enabled capabilities.* U.S. DoD. https://www.ai.mil/Portals/137/Documents/Resources%20Page/CDAO_TE_Framework_-_OTE_TES_2024-04-compressed.pdf

Defense Acquisition University. (2025). *Program management fundamentals handbook.* https://www.dau.edu/sites/default/files/2025-01/DAU%20Project%20-%20Program%20Management%20Fundamentals%20Handbook.pdf

Department of the Air Force & MIT. (2022). *Artificial intelligence acquisition guidebook.* https://atarc.org/wp-content/uploads/2022/11/Air-Force-MIT-Guidboik..pdf

Developmental Test, Evaluation, and Assessment. (n.d.). *Test & evaluation of artificial intelligence enabled systems (AIES).* https://www.cto.mil/dtea/te_aies/

Director, Operational Test & Evaluation. (2024). *Operational test and evaluation and live fire test and evaluation of artificial intelligence-enabled and autonomous systems* (DoDM 5000.101). U.S. DoD. https://www.dote.osd.mil

Gamboa, E. (2021, November 17). *NAVWAR announces Project Overmatch prize challenge winners.* U.S. Navy. https://www.navy.mil/Press-Office/News-Stories/Article/2849143/

Hicks, K. (2022, June). *DoD responsible artificial intelligence strategy and implementation pathway.* U.S. DoD. https://media.defense.gov/2024/Oct/26/2003571790/-1/-1/0/2024-06-RAI-STRATEGY-IMPLEMENTATION-PATHWAY.PDF

National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119–60 (2025). https://www.congress.gov/bill/119th-congress/senate-bill/1071

National Institute of Standards and Technology. (2025). *Adversarial machine learning: A taxonomy and terminology of attacks and mitigations* (NIST AI 100-2e2025). https://csrc.nist.gov/pubs/ai/100/2/e2025/final

National Security Commission on Artificial Intelligence. (2021, March). *Final report.* https://reports.nscai.gov/final-report/

PEO IWS. (2022, March 30). *ICS engagement brief.* NAVSEA. https://www.navsea.navy.mil/Portals/103/Documents/Exhibits/SAS2022/1100_ICS%20Engagement%20brief%2003302022_SAS_Distro%20A.pdf

Secretary of War. (2026, January 9). *Artificial intelligence strategy for the Department of War* [Memorandum]. U.S. DoW. https://media.defense.gov/2026/Jan/12/2003855671/-1/-1/0/ARTIFICIAL-INTELLIGENCE-STRATEGY-FOR-THE-DEPARTMENT-OF-WAR.PDF

State Council of the People's Republic of China. (2017, July). *New generation artificial intelligence development plan.* https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/

Toh, A., & Gledhill, J. (2024). *How acquisition reform could make military AI more expensive and less safe.* Brennan Center for Justice. https://www.brennancenter.org/our-work/analysis-opinion/how-acquisition-reform-could-make-military-ai-more-expensive-and-less

U.S. DoD. (2020a). *Operation of the adaptive acquisition framework* (DoD Instruction 5000.02). https://www.dau.edu/dod-instruction-500002-operation-adaptive-acquisition-framework

U.S. DoD. (2020b, February). *DoD ethical AI principles.* https://www.ai.mil/Initiatives/Responsible-AI/

U.S. DoD. (2022). *2022 National Defense Strategy.* https://www.defense.gov/National-Defense-Strategy/

U.S. DoD. (2024). *Chief digital and artificial intelligence officer (CDAO)* (DoD Directive 5105.89). https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/510589p.PDF

U.S. GAO. (2023, June). *Artificial intelligence* (GAO-23-105850). https://www.gao.gov/assets/gao-23-105850.pdf

U.S. Navy. (2021, June 15). *NAVWAR launches second Project Overmatch prize challenge.* https://www.navy.mil/Press-Office/News-Stories/Article/2659328/

Vincent, B. (2023, January 10). *Navy's Task Force 59 reaches full operational capability.* DefenseScoop. https://defensescoop.com/2023/01/10/navys-task-force-59-reaches-full-operational-capability/


Acquisition Research Program
Department of Acquisition, Finance, and Manpower
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net