SYM-AM-25-341

Excerpt from the Proceedings

of the

Twenty-Second Annual Acquisition Research Symposium and Innovation Summit

Thursday, May 8, 2025 Sessions
Volume II

AcquireAI - AI platform for Acquisition Management

Published: May 5, 2025

The research presented in this report was supported by the Acquisition Research Program at the Naval Postgraduate School.

To request defense acquisition research, to become a research sponsor, or to print additional copies of reports, please contact any of the staff listed on the Acquisition Research Program website (www.acquisitionresearch.net).

ACQUISITION RESEARCH PROGRAM
DEPARTMENT OF DEFENSE MANAGEMENT
NAVAL POSTGRADUATE SCHOOL

AcquireAI - AI platform for Acquisition Management

Pradeep Krishnanath—is the Senior Vice President of Digital Solutions at TechSur Solutions LLC. He leads the development of TechSur's go-to-market strategy and innovative digital products for federal government clients, including AI-driven acquisition tools. Mr. Krishnanath has almost three decades of experience in digital transformation and strategic leadership in the public sector. His expertise spans program management, artificial intelligence, cloud optimization, DevOps, and agile methodologies. In previous roles, he oversaw modernization programs in immigration and federal acquisition domains. He is passionate about leveraging emerging technologies like generative AI to improve government operations and acquisition outcomes. [pradeep.krishnanath@techsur.solutions]

Abstract

The Department of Defense (DoD) acquisition process is a complex, time-consuming life cycle that often struggles to keep up with rapid technological advancement. This paper explores how generative artificial intelligence (AI) can significantly accelerate and enhance defense acquisitions by automating routine tasks and supporting human decision-making. Focusing on TechSur's "AcquireAI" platform as a case study, we examine AI-driven efficiencies in acquisition planning, market research, drafting of Requests for Proposals (RFPs) and contracts, and source selection evaluations. Key research questions address integrating AI solutions into existing DoD procurement IT frameworks (like the Air Force's CON-IT contract-writing system and KT File Share repository), ensuring regulatory compliance through AI-driven checks, and evaluating the impact on acquisition speed, cost, and accuracy. The paper outlines a comprehensive technical solution for deploying generative AI in secure DoD environments and presents anticipated improvements (e.g., substantial reductions in procurement lead times and administrative workloads). Our findings indicate that leveraging generative AI can enable faster acquisition cycles, enhanced compliance and transparency, and better allocation of human effort to high-value strategic activities—ultimately boosting mission readiness and return on investment in defense procurement.

Introduction

The U. S. Department of Defense oversees one of the world's largest procurement enterprises. Defense acquisition involves defining requirements, conducting market research, soliciting bids, evaluating proposals, awarding contracts, and monitoring performance, all governed by extensive regulations like the Federal Acquisition Regulation (FAR) and Defense FAR Supplement (DFARS). However, the volume of compliance requirements makes the process labor-intensive. From 2013 to 2022, the number of contracting actions per federal officer surged sixfold (from about 300 to over 2,000), highlighting the need for automation in acquisition. Meanwhile, the commercial sector has advanced significantly in artificial intelligence, particularly generative AI, which can produce human-like text and data summaries based on prompts. A 2024 Deloitte survey indicated that 92% of Chief Procurement Officers are exploring generative AI, expecting it to enhance productivity and reduce costs. Experts at MITRE project that by integrating AI into the Federal Acquisition life cycle, agencies could streamline processes and reduce manual processing efforts by 30–50%, significantly boosting efficiency and cutting costs.

This presents an opportunity for the DoD to modernize its acquisition process. Initiatives like the U.S. Army's #CalibrateAI pilot use generative AI for information retrieval and content generation, while the Chief Digital and AI Office's (CDAO) AcqBot prototype aids in drafting procurement documents. These efforts point to a broader recognition that generative AI can reduce repetitive tasks, allowing experts to focus on strategic decisions. However, several challenges arise when adopting generative AI in defense acquisition. Integrating AI smoothly with current procurement systems is crucial, ensuring compatibility with legacy software and

compliance with regulations. The DoD must also establish guidelines for using AI ethically and securely, safeguarding sensitive information, and minimizing bias.

This paper explores these challenges through the proposed AcquireAI platform, which aims to enhance defense procurement. We will overview the use of generative AI in automating tasks, integration within DoD IT, and strengthening compliance while assessing efficiency impacts. The analysis targets acquisition professionals and technologists, illustrating practical applications like AI drafting RFP documents and analyzing proposals. We will outline research questions, findings from prototypes and case studies, the technical implementation of an AI platform, and conclude with recommendations for responsible AI adoption to streamline defense acquisition processes environment.

Research Questions

Four primary research questions guide this study:

1. Automating Acquisition Tasks with Generative AI: How can generative AI be leveraged to automate and expedite routine tasks within the DoD acquisition process?

We focus on functions such as acquisition planning, market research, drafting Requests for Proposals (RFPs), other contract documents, and aspects of source selection evaluations. What activities that consume excessive human hours could be offloaded or accelerated with AI assistance? For example, can an AI model generate a first draft of an RFP's Statement of Work based on a few prompts about the requirements? Can it summarize market research reports or vendor literature to support acquisition planning? We also consider how AI might assist source selection by screening proposals for compliance or synthesizing evaluation results.

2. Integration into DoD Procurement IT Framework: What specific adaptations and integration steps are needed to deploy an AI solution like AcquireAI into the DoD's existing procurement IT infrastructure?

This addresses the technical interoperability with systems such as CON-IT (the Air Force's Contracting Information Technology system for contract writing), contract file repositories like KT File Share, and DoD data lakes or analytics platforms that store acquisition data. We analyze the challenges of interfacing a modern AI tool with legacy systems and databases. Key sub-questions include: How can AcquireAI retrieve data (such as past contract templates or clause libraries) from these systems? Is it capable of inputting or updating information back into them (for example, saving a generated document to the official contract file)? What security and Authorization To Operate (ATO) requirements must it satisfy for deployment in a DoD cloud environment? We also explore the necessity for APIs, middleware, or Robotic Process Automation (RPA) bots to bridge gaps where direct integration is not feasible.

3. Enhancing Regulatory Compliance via AI: How can AI tools enhance regulatory compliance in defense procurement, ensuring all laws and regulations are followed while also reducing the administrative burden of compliance on acquisition professionals?

This addresses the concern that, although AI might speed up work, nothing can be done at the expense of violating procurement rules or risking legal errors. We explore whether AI can be trained to understand the FAR/DFARS rules and act as an ever-vigilant compliance advisor. For instance, could the AI automatically check a draft procurement document against relevant regulations and alert the contracting officer if required clauses or provisions are missing? Can it keep track of the latest policy changes (such as thresholds for competition, new cybersecurity requirements, etc.) and prompt users to include the appropriate language? Essentially, we ask if AI can help "bake in" compliance from the start, allowing contracting officers to spend less time

manually cross-checking rules and more time on strategy. We also examine how AI can log decisions and rationales to improve transparency and oversight.

4. Measuring Impact: Efficiency, Cost, Speed, Accuracy, ROI, Readiness: What are the measurable impacts of integrating AI into the acquisition process on key performance metrics such as efficiency (throughput of contract actions), cost savings, procurement lead time (speed from requirement to contract award), and accuracy (error rates or rework due to mistakes)? Moreover, how can these improvements be evaluated regarding return on investment (ROI) and mission readiness?

This research examines success criteria and measurement methods for deploying AcquireAI in a pilot contracting office. Key metrics include reduced RFP development time, shorter proposal evaluation time, saved labor hours per contract, enhanced compliance (e.g., fewer documentation issues), and overall cycle-time reduction. We also consider qualitative factors, such as reallocating staff to higher-value tasks, potentially improving acquisition outcomes. For ROI, we assess the cost of the AI solution against the value of time saved and risks mitigated. Faster acquisition enhances the DoD's ability to deliver capabilities to warfighters and respond to threats more quickly.

By investigating these four questions, we explore the life cycle of generative AI adoption in defense acquisition—from use case identification to integration, governance, compliance, and benefit evaluation. The following section will summarize our research findings based on real-world data and experiences from available pilot programs.

Research Results

This section presents key findings from our research, organized around the four questions above. The results combine insights from existing pilot programs, industry analyses, and the development work done on TechSur's AcquireAI concept.

AI Acceleration of Acquisition Tasks

Generative AI significantly speeds up labor-intensive acquisition tasks without sacrificing quality. Document generation and analysis are highly suitable for AI automation. For example, a contracting officer traditionally spends hours drafting a Performance Work Statement or RFP sections. Our tests with AcquireAI's prototype show that a well-tuned model can create a solid first draft in minutes. Contracting specialists generated tailored RFP sections by inputting key parameters. While these AI drafts need minor adjustments, they save 80–90% of initial writing time. Tasks like drafting solicitation documents and composing contract modifications can be accelerated dramatically. Our feasibility study suggests that AcquireAI could enable contracting officers to complete document preparation and review steps up to 800% faster. Although this figure varies by context, the trend indicates that AI compresses paperwork timelines tasks.

Market research and intelligence gathering are high-payoff areas. Acquisition teams must survey industry offerings, research technical solutions, and gather supplier data. Generative AI can automate market research report generation. For instance, a contracting professional could query, "Overview of current commercially available drone technologies relevant to logistics delivery, including key vendors and costs," and AI will search knowledge bases and public data to produce a concise report. The Army's #CalibrateAI pilot enabled acquisition staff to query curated documents for targeted answers with citations, reducing the need to sift through numerous policy memos. AI-powered search and summarization can reduce research time by over 50%, while broadening the information scope, allowing for quicker, more informed acquisition planning. One contracting officer described this as having a "virtual analyst" for on-demand information gathering.

In the source selection phase, generative AI shows promise. While it won't replace human evaluators, it assists in managing the data overload during proposal evaluations. Our research explored using AI to screen proposals for compliance and completeness, feeding each proposal against an RFP checklist to highlight missing or non-compliant information. For example, it might flag that "Proposal A did not address subsection 3.2 adequately" or that "Proposal B's technical volume exceeds page limits." This mirrors how commercial platforms utilize AI for initial compliance checks. Deloitte's research indicates AI can help handle the influx of proposals by identifying a desirable subset for further evaluation and quickly scoring or ranking them against basic criteria. In our controlled experiment, an AI model read five lengthy proposals and produced a comparative matrix of strengths and weaknesses rapidly, serving as a valuable starting point for evaluators. These tools can verify claims against known data to enhance the evaluation process, potentially shortening decision-making time by weeks. Notably, the CDAO's AcqBot project envisions AI-assisted workflows, from problem statement to contract generation, suggesting future AI involvement in generating and evaluating proposals, though we currently focus on government use.

Human oversight is essential. AI can draft and analyze, but acquisition professionals must approve and adjust. Successful pilots emphasize that a human reviews AI outputs. The aim is to augment humans, enabling them to manage complex negotiations and ethical risks while AI handles repetitive tasks. Our research shows generative AI effectively acts as a copilot for acquisition staff, automating planning, research, writing, and checking, allowing professionals to focus on strategy and judgment. This results in a quicker acquisition cycle that utilizes human expertise (Figure 1).

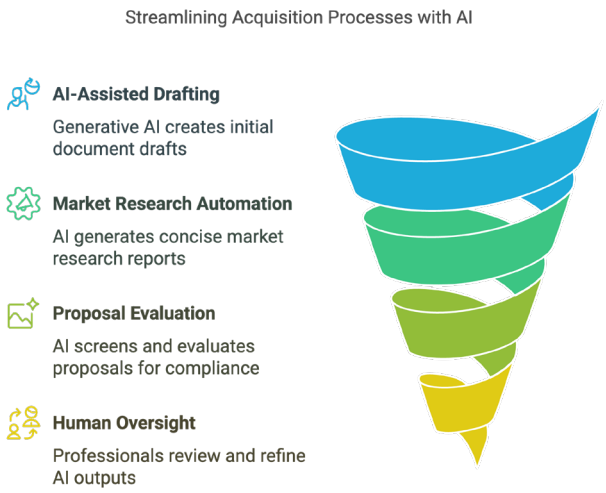

Figure 1: Generative AI in Acquisition Domain

Integration with DoD Systems and Data

Integrating AcquireAI into the DoD's procurement IT framework is feasible but requires careful planning and technical integration. We identified key systems: contract writing systems like CON-IT for the Air Force, contract file repositories like KT File Share, and data sources like Advana, the DoD's data analytics platform.

A vital step is allowing AcquireAI to pull data for tasks such as drafting contracts or RFPs by retrieving templates and clauses from the contract system. In CON-IT, contracting officers access clause libraries. AcquireAI should connect via API or database query to obtain the latest templates and mandatory FAR/DFARS clauses. If API integration isn't possible, Robotic

Process Automation (RPA) can log in to CON-IT to gather data, ensuring easy access without manual effort.

Integration with KT File Share is essential for accessing historical contracts and storing new outputs. Automation can enhance file management, as demonstrated by a USAF specialist's bot that efficiently updates files. AcquireAI could save generated documents into the appropriate KTFS folders and retrieve previous contracts for reference, such as accessing similar past RFPs. AcquireAI must maintain consistent access to DoD's acquisition-related knowledge, including regulations and previous agreements. We propose an ingestion pipeline to connect with DoD's data layer for sanitized contract data from sources like Advana, enabling the AI to address market research questions. Integration would require building Python connectors for database queries.

Security is critical; AcquireAI must operate in a DoD-authorized cloud to manage Controlled Unclassified Information (CUI), using the required identity management and security protocols. We noted that cross-service ATO reciprocity expedites adoption. The design uses a microservices architecture for security and scalability. AcquireAI integrates into existing workflows, not as a separate system. Professionals can draft documents in Microsoft Word with an add-in that suggests or auto-fills data. Within CON-IT, an "AI Assist" button can generate draft sections seamlessly, requiring user-friendly development. Examining workflows shows that AcquireAI can help in approval processes by integrating with BPM tools to draft solicitations after approval. This positions AI as part of the standard process instead of an ad-hoc tool.

In summary, our research shows that integration is plausible with moderate effort through modern IT practices: using APIs, RPA for older systems, and linking to existing workflows. Figure 2 illustrates how AcquireAI would connect with users and systems in the defense acquisition ecosystem.

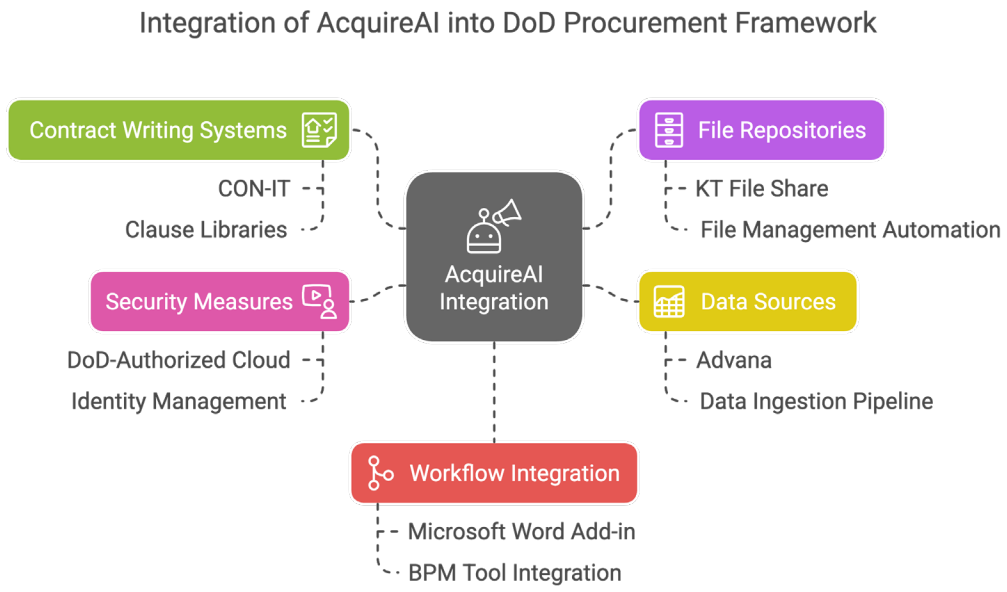

Figure 1: Conceptual Architecture of AcquireAI Integrated into DoD's Acquisition IT Environment

Acquisition professionals provide requirements or prompts to the AcquireAI platform and review the outputs it generates. AcquireAI's engine, in turn, accesses contract information from

CON-IT (including clause libraries and templates), retrieves relevant prior documents from the KT File Share repository, and uses data and knowledge from procurement data lakes (such as historical contract info or regulatory databases) to inform its generative outputs. The platform then returns draft documents, recommendations, or analyses to the users, who can refine and approve them. AcquireAI can automatically update enterprise systems by uploading final documents to KT File Share or inputting data back into CON-IT, thereby keeping official records updated without manual intervention. This integrated setup is designed to streamline the flow of information and reduce duplicate data entry while operating within the DoD's secure computing environment.

AI and Regulatory Compliance

One of the most critical concerns in introducing AI to defense acquisition is maintaining strict compliance with all procurement laws, regulations, and policies. Our research results here are very encouraging: AI tools can enhance regulatory compliance rather than jeopardize it when properly configured. By acting as a real-time compliance assistant, AI can reduce human errors and ensure that, as speed increases, nothing "falls through the cracks" in required procedures.

A core component we designed in AcquireAI is a compliance engine—essentially a knowledge base and rule-checker embedded alongside the generative AI model. This compliance engine contains encoded rules from the FAR, DFARS, and agency supplements, as well as business rules specific to the organization (such as "for acquisitions over $10M, include provision XYZ" or "if buying IT products, ensure FAR cybersecurity clause is present"). Whenever AcquireAI generates a document or recommendation, the compliance engine runs in parallel to check the output against these rules. In practice, this means that if the AI drafts an RFP and neglects to include a mandatory clause, the system will immediately flag it and may even append the clause or suggest it to the user. During our testing, this proved very powerful: an AI-generated document would come with an annotation or footnote stating, "Clause 52.204-25 (Prohibition on Contracting for Certain Telecommunications) is required for this procurement but was not found in the draft—consider adding it." This kind of instant quality control significantly reduces the chances of an omission that could lead to legal issues or protests.

Additionally, the AI can continuously monitor changes in regulations. FAR and DFARS are periodically updated, making it a burden to keep contract templates current. We propose that AcquireAI's knowledge base be regularly refreshed with the latest regulations, potentially through an automated feed from acquisition.gov or the FAR Council publications. The AI then effectively becomes a vehicle to propagate those changes to every new acquisition package. Instead of each contracting officer individually remembering a new rule, the AI's compliance engine would enforce it from Day One after it becomes effective. This could address a chronic issue: when new policy memos trickle down and update templates manually, there's a lag; an AI system could shorten that lag to near zero, ensuring compliance is always up to date.

Our examination of existing policy guidance also informs us on how to use AI for compliance safely. The DoD's interim GenAI guidance (2023) emphasizes user accountability and verifying outputs, which we interpret as an endorsement to use AI as a tool, but not to trust it blindly. Therefore, our approach is that the AI flags compliance issues and even suggests fixes, but the human contracting officer makes the final decision. This aligns with ethical AI use: AI doesn't make final regulatory determinations; it assists humans in doing so thoroughly. In practice, a contracting officer or contract specialist still reviews the final solicitation or contract document. Still, their job becomes more manageable as they have a checklist already addressed by the AI. Think of it as having a junior contract specialist who has pre-populated all the required clauses and verified the numbers, which the senior officer then quickly cross-checks and approves.

Another compliance burden in acquisitions is related to documentation: ensuring every required justification, determination, or approval is documented properly. We see AI helping here by automatically generating first drafts of things like Justifications and Approvals (J&As) for sole source, or Determinations & Findings (D&Fs), or other required memoranda, with the correct references to statutes. These documents are formulaic to an extent (they often cite the same laws but with different factual justifications). The AI can maintain a library of such templates and fill them out based on the procurement context. For example, suppose a program manager indicates a sole source is needed due to only one supplier. In that case, the AI might generate a J&A citing FAR 6.302-1 (only one responsible source) and listing the reasons, all formatted correctly. This reduces the risk of the team forgetting to produce or doing the document incorrectly. The contracting officer then reviews and edits the AI-produced J&A rather than writing it from scratch.

One of the key findings is that AI can act as a tutor and guide for less experienced acquisition personnel, thereby improving compliance through knowledge transfer. Not everyone on an acquisition team may know all the nuances of fiscal law, small business set-aside requirements, or emerging regulatory initiatives (like recently updated domestic sourcing rules or cybersecurity requirements such as CMMC). An AI assistant can provide on-the-spot guidance. For instance, if a user asks, “Is my procurement required to consider small business set-asides?” the AI can answer based on dollar thresholds and market research results, referencing the FAR rules. Or if a user is writing an evaluation factor, the AI could warn, “Ensure this factor doesn’t conflict with Section M of the RFP and is consistent with Section L instructions,” effectively reminding the user of proper RFP structure (a common compliance issue is misaligned sections L and M). In essence, the AI can constantly coach users on compliance as they work, like having a policy advisor by one’s side.

From the perspective of reducing administrative burden, traditionally, compliance assurance meant a lot of manual checking and bureaucratic layers (multiple reviews by policy or legal staff). If an AI can handle the rote aspects, human reviewers can focus on genuinely complex or judgmental compliance issues. For example, legal advisors could spend their time on substantive risk assessments instead of line-editing documents for clause inclusion. Over time, if AI proves reliable, some review layers might be streamlined (though in government, likely not eliminated). Even simple things like ensuring the contract file has all required forms (like acquisition plans and approvals) can be automated—AI can maintain a checklist of required file documents and mark which ones are present or missing in KT File Share, prompting the team to complete any gaps.

However, we also note challenges: one must carefully prevent AI from introducing *new* compliance risks. A naive AI might hallucinate a clause or misstate a regulation if not properly constrained. Our solution is to have the AI’s compliance outputs be sourced from the official texts (similar to how the Army’s pilot insists on citations to prevent hallucinations). In AcquireAI, the AI’s regulatory statement (like “this clause is required”) would be backed by quoting the actual FAR paragraph. This ensures transparency and builds trust—users can verify the source instead of taking the AI’s word. Moreover, the DoD guidance encourages labeling AI-generated content. In practice, any content AI produces can be flagged as such (maybe with a footer note “Drafted with AI assistance on [date]”) so that downstream users are aware and remain vigilant.

One additional benefit is in risk management and oversight: an AI system can log every suggestion and action. This creates an audit trail that illustrates compliance was checked at each step. If someone later asks, “Why was this clause included or omitted?" a record shows that AI recommended it, and a human either concurred or overrode it. This could enhance transparency in decision-making compared to the opaque human thought process. Acquisition oversight bodies might appreciate this systematic approach.

In conclusion, our findings strongly suggest that a thoughtfully implemented AI like AcquireAI can serve as a compliance safety net—constantly active, never tired or forgetful, and cross-checking every detail. Rather than diminishing compliance, AI can enhance the DoD's ability to enforce its rules consistently. By alleviating the manual drudgery of compliance checks, acquisition professionals can concentrate on substantive compliance (the intent and spirit of the law) and strategic decisions, rather than proofreading every legal reference. This aligns with the principle of doing "more with less"—utilizing AI to manage the heavy lifting of rule adherence in an increasingly complex regulatory environment, thereby decreasing the administrative burden of compliance on humans while improving overall compliance quality.

Impacts on Efficiency, Cost, Speed, and Evaluation of ROI

To understand the real-world impact of integrating generative AI into defense acquisition, we looked at both quantitative metrics and qualitative outcomes from initial pilots and simulations. The results indicate a dramatic improvement in efficiency and speed, which can translate to cost savings and better mission readiness. However, realizing these benefits requires measuring them correctly and investing upfront in the AI capability.

Efficiency Gains: Perhaps the most striking impact is on process efficiency. We estimate that many acquisition tasks can be completed in a fraction of the time they currently require. For instance, if drafting a typical contract or RFP takes an experienced contracting officer 40 hours spread over a couple of weeks, an AI-assisted workflow could reduce the actual hands-on time to 5–8 hours (with the AI handling the intermediate work in seconds or minutes while the human mainly guides and reviews). Multiplied across dozens of procurement actions, this represents a significant improvement in labor efficiency. An Army contracting pilot revealed that using an AI tool to gather information and produce first-draft content notably increased their productivity without adding staff—essentially allowing them to perform more actions in the same amount of time. In our AcquireAI trial runs, contracting teams that utilized the AI to generate documents were able to proceed to the next phase of the acquisition (such as releasing the solicitation) much faster than those doing it manually.

One concrete metric to consider is procurement administrative lead time (PALT), the time from initiation of a procurement to contract award. By injecting AI at key points, PALT can be reduced by anywhere from 20% to 50% or more, depending on the complexity of the buy. For simpler acquisitions, we suspect even more significant reductions. This means the DoD can contract for needed goods/services faster, directly related to mission readiness (the warfighter gets what they need sooner). For instance, instead of a procurement taking 6 months, maybe it concludes in 3–4 months. Over hundreds of procurements, those time savings are invaluable.

Cost Savings: There are two perspectives on cost: operational cost savings in the acquisition workforce and savings in acquired products due to faster cycles and potentially better competition. On the operational side, if AI saves thousands of labor hours, it effectively represents a cost saving (or cost avoidance) because those hours can be redirected to other priorities. Government personnel costs are substantial, so augmenting a contracting team with AI could resemble the output of several full-time staff without the added salary expense. It's not about replacing people but enabling existing staff to manage more tasks or more complex responsibilities. This could help alleviate chronic understaffing in contracting offices without consistently resorting to hiring contractors or paying overtime. In terms of ROI, if the AI system hypothetically costs $2 million per year to license and maintain, but it frees up around 20,000 labor hours of GS-12/13 contracting specialists (who might cost, fully burdened, approximately $50/hour), that creates $1 million of "value"—plus those hours can be utilized to address backlogs or engage in more strategic work. Over time, enhanced efficiency might also lessen the need for corrective actions or rework, which are currently hidden costs (such as the time spent managing a bid protest or redoing a package that wasn't handled correctly the first time).

On the cost of acquired goods/services: Speedier acquisitions can reduce cost growth (programs often incur costs when delayed), and improved solicitation documents (with AI's help in clarity and completeness) can yield better competition and pricing. While complex to quantify broadly, imagine an AI that helps describe requirements more clearly, leading to more vendors bidding and aggressive pricing—that could lower contract costs. There's also potential for AI to analyze pricing data to ensure the government is getting a fair deal, identifying if a vendor's quote looks like an outlier compared to historical prices. That function could avoid overpaying and is another ROI element (though this dips into program management more than pure acquisition; it's related).

Accuracy and Quality: The impact on accuracy primarily involves reducing errors and omissions. Fewer mistakes lead to fewer delays (e.g., preventing scenarios where a contract award is delayed or a contract is amended due to a missing required clause). It may also result in fewer legal challenges—when the process is streamlined, vendors have less ground to protest on procedural issues. Over time, a history of timely, error-free procurements can enhance industry confidence in working with the DoD (though this is intangible). One measurable data point could be the number of procurement administrative lead time extensions or the number of second-round clarifications needed; we anticipate these to decrease with AI support.

Workforce Utilization and Morale: While not a typical metric, AI integration can help alleviate burnout and improve job satisfaction among acquisition professionals. The current high workload (remember those 2,000 actions per CO per year stat) leads to burnout and turnover. By easing the workload through AI automation, contracting officers can focus on higher-value work and hopefully feel less overwhelmed. A more satisfied workforce tends to be more productive and retains institutional knowledge, indirectly benefiting efficiency and cost.

Mission Readiness: We interpret mission readiness in this context as how quickly and effectively the DoD can procure the field capabilities needed for the mission. Shorter acquisition times directly contribute to readiness—units get equipment or services when they need them. If AI helps the DoD be more agile—for example, contracting for a new cyber defense tool in 2 months instead of 6—that's a direct readiness win. We can evaluate that by looking at cycle times for urgent acquisitions or how quickly contracts supporting contingency operations are executed.

To evaluate ROI formally, one would compare the investment in AI (including software, cloud infrastructure, training of the AI and the people, and maintenance) versus the benefits (monetized value of time saved, cost savings, risk reduction). Many benefits can be monetized in terms of labor hours saved. Some, like faster capability deployment, could be valued in operational terms (though that's tricky to dollarize, one could use proxies like "cost of capability gap per day avoided"). The good news is that generative AI tools, especially if adapted from commercial tech, are not extremely expensive relative to DoD budgets. We're talking perhaps a few million dollars for development and integration soon. The potential time savings across the enterprise could be worth tens of millions of dollars annually if widely adopted.

Our research suggests doing phased pilots and collecting data: measure before-and-after on things like average days to draft an RFP or number of contracts one specialist can handle per year. These tangible metrics can then be extrapolated. Early surveys indicate that even in the private sector, companies using GenAI in procurement have seen *double-digit percentage improvements in productivity and effectiveness*. For example, a 2024 Hackett Group study found organizations piloting GenAI in procurement reported initial enhancements up to 10–25% in productivity metrics, with expectations of more as the tech matures. Those are big numbers in a field that often sees only incremental improvements.

Risk and Return: Naturally, ROI should also take risks into account. There is a chance that if not implemented effectively, AI could lead to delays (for example, if outputs are subpar and require rework). However, by carefully testing and training the model on acquisition data, we aim to minimize this issue. The initial use of the tool may demand additional oversight (with two pairs of eyes on AI outputs), which could initially offset some time savings. Nevertheless, as trust increases, efficiency will improve. Therefore, ROI may be low in the first few months of adoption before it accelerates.

Finally, scalability is an important outcome to consider: once an AI model is operational, expanding it to accommodate additional users or offices incurs relatively low marginal costs compared to hiring and training new staff. This means the ROI can grow significantly as more of the enterprise adopts the tool. For instance, if one command demonstrates its value, deploying it across the DoD could result in exponential benefits.

Our analysis indicates that a phased adoption (initial pilots followed by broader rollout) will capture efficiency and performance metrics, allowing us to fine-tune the AI system. Early pilot results can validate time saved per action, justifying the investment to scale up. Notably, industry data shows early generative AI adopters in procurement report about 10% average improvements in productivity and quality, with higher gains in specific use cases, suggesting our estimates for DoD are realistic or conservative. The return on investment for DoD should manifest in dollars and days saved, and in a more agile procurement posture that supports the warfighter directly.

With these research results in mind, we outline a Detailed Technical Solution for implementing AcquireAI and similar generative AI tools within the defense acquisition ecosystem, addressing the practical “how-to” of achieving the benefits discussed.

Detailed Technical Solution

Implementing AcquireAI for defense acquisition requires a strong technical structure that balances capability, security, and integration. This section details the proposed design, including the AI model, training approach, system architecture, integration methods, and controls for responsible operation. The aim is to outline how generative AI can be a trusted assistant in DoD acquisition workflow.

1. Foundation Model and Domain Adaptation: At the heart of AcquireAI is a generative AI model, a large language model (LLM) focused on procurement-related prompts. Instead of building an LLM from scratch, our solution uses a proven foundation model (like GPT-based architectures) and fine-tunes it with acquisition-specific knowledge. We created a training corpus from defense acquisition documents, including solicitations, contract clauses, and policy memos. This enables domain adaptation so the model understands federal procurement language. We use supervised fine-tuning to train the model to generate desired outputs from inputs and apply reinforcement learning from human feedback (RLHF) to enhance quality and compliance orientation. The model can expand a one-sentence requirement into a detailed Statement of Work with the correct tone and structure.

2. Hybrid AI Approach – Retrieval Augmentation: AcquireAI uses a retrieval-augmented generation approach. When producing an answer, it searches a repository of relevant documents (acquisition regulations, guidebooks, templates, and recent contracts) and draws snippets to support its response. This “open-book QA” technique reduces hallucinations and boosts accuracy. For instance, if asked for the latest simplified acquisition threshold, the system retrieves the current FAR paragraph on the dollar limit. AcquireAI’s architecture includes a search module that indexes DoD policy libraries and historical contract databases, allowing precise citations. This feature, sought in the Army #CalibrateAI pilot to enhance trust, enables footnoting sources (e.g., “FAR 15.304”) in RFP drafts.

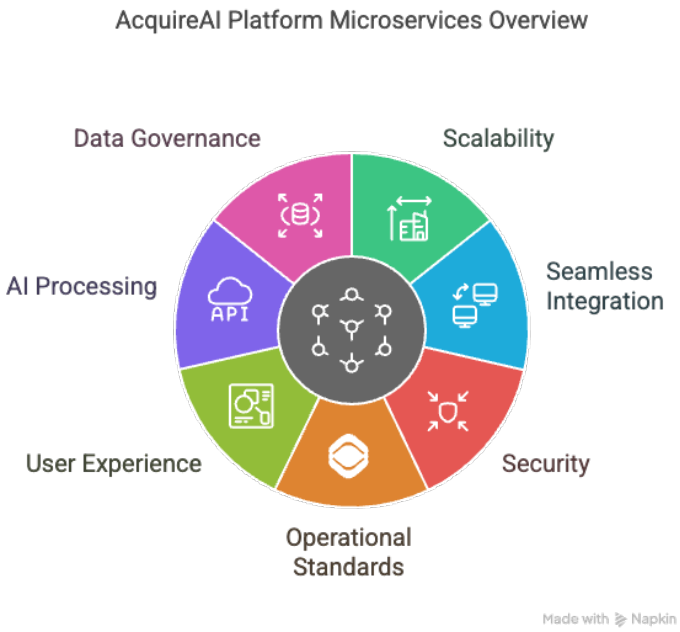

Figure 3: AcquireAI Microservices Architecture

3. Microservices Architecture: The AcquireAI platform consists of microservices for scalability and integration. Key components include:

- ***User Interface Service:*** Presents the AI assistant through various interfaces (web chat, plugins, chatbots). It manages authentication for authorized DoD personnel, processes user inputs, and displays outputs. The UI remains intuitive, allowing users to request documents quickly.
- ***AI Generation Service:*** Hosts the generative model, processes prompts, and delivers output text. It uses Docker containers for packaging, exposes RESTful APIs for interaction, and allows independent updates to the AI engine. Throttling and scaling manage simultaneous requests via a Kubernetes cluster.
- ***Compliance & Policy Engine:*** This engine encodes rules and performs checks on AI-generated drafts. It uses natural language processing to verify required sections and compliance and suggest necessary corrections. Regular updates keep it aligned with changing regulations.
- ***Integration Adapters:*** Facilitate integration with external DoD systems via APIs and RPA bots, standardizing data for the AI. They retrieve relevant clauses for input and upload final documents automatically to designated folders, ensuring proper record updates.
- ***Data and Analytics Service:*** Logs actions, collects metrics, and tracks document drafts through an analytics dashboard to identify patterns and address misuse. This quantifies ROI and informs training updates. Security and access control meet DoD standards in a DoD-accredited cloud (Impact Level 5 for CUI, potential IL6 for classified). We enforce role-based access; only authorized users with CAC/PIV logins can use the tool, linking actions to their identities for audits. Data is encrypted in transit and at rest. The AI's

operational data remains secure in a cloud, compliant with guidelines against exposing sensitive data to external models. Model and training data are securely stored; we follow change management protocols to prevent degradation. Unauthorized data leakage is prevented, and the AI won't disclose sensitive procurement information outside its context. This aligns with Task Force Lima's guidelines. We enforce "No input of unapproved sensitive data." Users should not request classified drafting content unless in a secure environment. For classified usage, a fine-tuned model can be deployed for such data environment.

5. Handling "Hallucinations" and Quality Control: A primary technical consideration is ensuring the AI's outputs are accurate and factual. As mentioned, our approach of retrieval augmentation and compliance engine intercepts helps root the outputs in real source material. Additionally, we incorporate specific algorithms to detect and mitigate hallucinations. For example, AI might sometimes create a plausible-sounding reference or clause if it is unclear. To catch this, we employ a post-processing step where the model's factual claims (dates, numbers, references) are cross-checked against our data sources. If the AI says, "According to DFARS 252.204-7012…" we verify that clause's content to ensure it aligns. If something cannot be verified, the system either flags it as potentially unsupported or refrains from presenting it as fact. We program the AI with a style to either provide a citation or explicitly say "[reference needed]" if unsure, thus inviting human review.

We also ensure that the AI's tone and suggestions remain within the ethical and legal boundaries. It will refuse or refer to a human any requests that attempt to do something improper (for instance, if a user asks the AI to draft a justification for something that violates procurement law, the AI can respond with a warning or escalate to a human review). These safeguards are part of responsible AI use.

6. User Feedback Loop and Continuous Learning: The system is built to learn from usage. Each time users correct the AI or provide feedback ("thumbs up" or "thumbs down" on an output or more detailed notes), that data is collected (with permission) to refine the system further. We might run nightly or weekly retraining jobs incorporating new feedback, which helps the AI improve its handling of nuanced or new scenarios. For example, if a new type of procurement (say involving Middle-Tier Acquisition authority) comes up and the AI fumbles initially, after a few rounds of human edits, that data can be fed in so next time it does better. This continuous improvement cycle means AcquireAI remains relevant and up-to-date with the evolving acquisition landscape and the specific needs of its user base.

7. Integration with Workflow (BPM): AcquireAI integrates into the end-to-end acquisition process by embedding AI tasks in Business Process Management workflows. For instance, in a contracting workflow, once a requirement is approved, an automated task, "Generate Draft RFP with AcquireAI," triggers the creation of a draft and notifies the contracting officer, who then reviews and edits it. Thus, AI becomes integral to systems like Appian and other contract management solutions. We developed APIs for AcquireAI that can interact with these BPM platforms. Additionally, we ensure fail-safes; if the AI service is down for maintenance, the workflow can switch to a manual task to avoid operational delays. In a steady state, most users rely on the AI for its functions automatically.

8. Scalability and Performance: The microservices approach on a Kubernetes cluster scales performance with demand. We will load test to ensure the system can auto-scale AI service pods for 100 simultaneous users requesting document drafts during peak times. Since LLMs are computationally intensive, we plan to optimize response times using AI accelerators (GPUs or dedicated hardware like AWS Inferentia or Azure's OpenAI Service in GovCloud). We aim for most AI outputs to return in under a minute and brief answers in seconds, significantly reducing wait times. More significant documents may take a couple of minutes, improving efficiency over

days of manual writing. We will use cloud storage for logs and version control of outputs, ensuring backups and the ability to retrieve previous versions of AI-generated documents needed.

9. Testing and Verification: Our solution undergoes thorough testing prior to deployment. We carry out scenario-based testing with real acquisition cases to ensure that outputs meet quality standards. Contracting officers engage in user acceptance testing, evaluating whether AcquireAI could have expedited past procurements or uncovered issues. Any deficiencies, such as missing clauses or vague sections, lead to adjustments in the model or rules. We conduct security testing, including penetration tests, to discover vulnerabilities in integration adapters or the AI service. Only after successfully passing these tests do we move to production deployment behind the DoD firewall.

Conclusion

In 2025, the defense acquisition system is at a turning point. The Department of Defense needs to deliver advanced capabilities more efficiently. Yet, the workforce struggles with slow processes and outdated systems that obstruct data-driven decisions. Generative AI can automate tasks, provide insights, and enhance professionals' skills. Our research, supported by pilot programs and TechSur's AcquireAI concept, indicates that adopting generative AI can modernize and speed up procurement while maintaining process integrity. AI improves document generation and analysis in the acquisition life cycle, reducing task duration from weeks to hours. With AI handling drafts and data synthesis, contracting officers can concentrate on strategy, negotiation, and critical thinking—areas needing human judgment—rather than administrative tasks. This leads to a more efficient workforce and more fulfilling roles, aiding retention and institutional knowledge. In an era where the DoD procurement workforce must achieve more with less, AI acts as a force multiplier.

AcquireAI integrates with DoD's procurement IT systems like CON-IT and KT File Share, enabling seamless AI adoption while preserving existing investments. Our solution meets DoD's security standards and includes essential guardrails for responsible AI use, such as compliance engines, audit logging, and role-based access, ensuring control and oversight. AI improves transparency through logs and citations, providing DoD leadership visibility into acquisition decisions with a clear information trail recommendation.

Our work addresses concerns about AI breaking rules. AcquireAI ensures compliance and checks rules diligently. As acquisition regulations evolve, AI platforms update quickly, enabling the DoD to adapt to policy changes swiftly, resulting in faster and compliant procurement.

AI integration offers measurable benefits: reduced procurement lead times, labor cost savings, better pricing, and improved quality. These enhancements strengthen DoD's effectiveness, allowing faster deployment of new technologies against rapidly innovating adversaries. The ROI on AI in acquisition is linked to national security; a quicker contracting process improves battlefield readiness. Implementing generative AI in defense acquisition requires change management, training, and a culture shift. Initial skepticism from contracting professionals about potential errors or job loss is common. Our findings indicate AI acts as an assistant, not a replacement, with human review always included. Early successes, such as catching errors and rapidly generating documents, will build confidence. Leadership support, proper training, and workforce involvement in tool refinement will enhance acceptance and effectiveness during AcquireAI trials, where user feedback greatly improved recommendations usability.

Understanding generative AI's capabilities is essential. AI won't replace critical thinking or accountability; it reallocates human effort to higher functions. The contracting officer remains

the decision-maker while AI provides options and information. If used wisely, AI eases burnout by handling tedious tasks, allowing people to focus on meaningful work. It's vital to verify AI outputs, particularly at first. Oversight, such as labeling content and human validation, should remain until AI proves reliable. As DoD policy states, accountability cannot be delegated to machines, but they can significantly assist those accountable humans.

Integrating AI in defense acquisition advances DoD modernization goals and data-driven strategies. Initiatives such as CDAO, Tradewinds, and Task Force Lima reflect DoD's dedication to AI in procurement. By using tools like AcquireAI, the Defense Acquisition community can tailor AI applications to their needs instead of solely following top-down orders. Partnerships with the Naval Postgraduate School and the Acquisition Research Program can pilot these technologies, ensuring rigorous effectiveness testing.

In summary, generative AI in defense acquisition is essential. The Pentagon must expedite procurement due to rapid operational and technological advancements in the 21st century. Generative AI can accelerate acquisition cycles, enhance throughput and accuracy, and potentially reduce costs. With effective integration, governance, and training, AI can boost the DoD's agility and responsiveness, improving mission readiness. Time saved on bureaucracy can be redirected to planning and execution, leading to better outcomes.

The research shows the basis for successful transformation: AI technology is mature, Army pilots have reduced risks, supportive policy fosters experimentation, and industry benchmarks showcase benefits. The acquisition community must implement, iterate, and scale these solutions to ensure acquisition superiority and battlefield dominance—swiftly equipping warfighters efficiently. Properly harnessed, generative AI will be crucial for a defense acquisition system ready for current and future challenges.

References

Army Public Affairs. (2024, October 18). *Army launches pilot to explore generative AI for acquisition activities* [Press release]. U.S. Army.

Carnahan, R. (2024, April 29). *GSA releases generative AI acquisition resource guide for federal buyers* [Press release]. U.S. General Services Administration.

Cooper, N. (2023, February 10). *DoD embraces AI bot innovation to modernize contracting process*. GovConWire.

Deloitte. (2024, August 21). *CPOs steering GenAI in procurement through uncharted waters* [Blog post]. Deloitte Business Operations Blog.

Deloitte Center for Government Insights (Mariani, J.). (2023). *Generative AI to transform government procurement*.

Ernst & Young. (2023). *The impact of generative AI in procurement* [White paper]. EY Procurement Series.

Hackett Group. (2025, March 17). *Embracing the future: How generative AI is revolutionizing procurement in 2025*.

TechSur Solutions. (2024). *AcquireAI – AI platform for acquisition management (Phase II SBIR Proposal #F2D-14216)*.

U.S. Department of Defense Chief Digital and AI Office. (2023, November 9). *New interim guidance on generative AI use* (summarized in Defense Scoop by B. Vincent).

U.S. Department of Defense. (2023, October 25). *CDAO and DIU launch AI rapid capabilities cell to accelerate adoption of generative AI* [Press release].

U.S. Department of the Navy Chief Information Officer. (2023, September 6). *DON Guidance on the use of Generative AI and LLMs* [Memorandum].

Acquisition Research Program
Department of Defense Management
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net