Accelerating and Improving Decision Making: A Vision for Agentic AI in U.S. Defense-Industrial Collaboration

I. Executive Summary

The United States' enduring advantage in national security has always been rooted in our capacity for innovation and our ability to mobilize the full strength of the defense industrial base (DIB). However, the speed and complexity of the modern world are straining our traditional decision-making frameworks. The volume of data, the fragility of global supply chains, and the pace of technological change often outstrip the capacity of human-centric processes.

Agentic artificial intelligence (AI) presents a profound opportunity to address these challenges. Unlike traditional AI, which typically performs a single, reactive task, Agentic AI refers to a new class of proactive, goal-oriented systems that can perceive an environment, plan a course of action, use relevant digital tools, and responsibly execute a plan autonomously or with human oversight to achieve a specific objective. For senior leaders in the Department of War (DOW) and the DIB, this is not merely an incremental improvement — it is a fundamental shift that can significantly reduce administrative burdens, accelerate timelines, and provide a new level of strategic foresight.

This paper outlines a vision for what is possible and proposes several sample use cases for Agentic AI for both the DOW and companies across the DIB. The ultimate goal is to build a more adaptive, resilient, and strategically superior defense ecosystem — one that is not simply human-powered, but human-augmented.

II. Introduction

The modern defense landscape is defined by an acceleration of complexity. From intricate global supply chains to the rapid evolution of technology, the sheer volume of information and the speed at which decisions must be made present an unprecedented challenge. The traditional, sequential, and often bureaucratic processes that have long governed the relationship between the DOW and the DIB are struggling to keep pace. This is not a failure of will or talent, but a limitation of scale and speed inherent in a human-driven system.

A new paradigm is required. This paper introduces Agentic AI as that paradigm. This new class of technology moves beyond simple automation to enable a system capable of independent action toward a stated goal. These are not static tools that await a command, but dynamic partners that can proactively analyze data, reason through complex situations, and work collaboratively with humans to solve problems.

This document serves as a strategic guide for senior decision-makers. It does not delve into the technical underpinnings of AI or machine learning (ML) algorithms. Instead, it will focus

on the tangible benefits of Agentic AI, its potential to revolutionize the partnership between the DOW and its industrial partners, and the practical considerations required to begin this journey.

III. Understanding the Agentic AI Advantage

To appreciate the strategic value of Agentic AI, it is essential to understand how it differs from the reactive, rule-based automation with which we are familiar. A traditional AI system provides directional guidance based on predefined rules. If an unexpected variable is introduced, such as a road closure, the system may simply prompt a re-evaluation or provide a new set of rigid instructions. Agentic AI, in contrast, is capable of a more dynamic approach. If it detects an unexpected event, it will autonomously replan the entire route, re-optimize for time and safety, and inform the operator of the new path and the reasoning behind its choice.

The core capabilities that enable this are:

Proactivity: Agents do not wait for a prompt; they initiate actions based on environmental triggers and defined objectives. For instance, instead of waiting for a human to request a report, an agent could monitor a project timeline and automatically generate a status update whenever a key milestone is at risk.

Autonomy: Within carefully defined boundaries and guardrails, agents can operate independently to achieve their goals. This frees human decision-makers from time-consuming, low-level tasks, allowing them to focus on high-level strategic challenges that require judgment and experience. Depending on the task's risk profile, there can be multiple options for human involvement: humans can be in-the-loop or on-the-loop, or the process could be fully autonomous with after the fact data auditing.

Collaboration: Agents are not isolated tools; they are designed to work together and with human users. A team of agents could, for example, collaborate to analyze disparate datasets and present a unified, actionable briefing to a human analyst. A supervisory agent could coordinate this team of agents, assign tasks, and evaluate outcomes.

Learning and Adaptability: An agent can observe the results of its own actions, reflect on its performance, and autonomously improve its plans and actions. This creates a powerful feedback loop that ensures the system becomes more effective and efficient over time.

IV. Strategic Applications in Defense-Industrial Collaboration

Application 1: Streamlined Acquisition and Procurement

The defense acquisition process is complex and time-consuming. From identifying a need to issuing a contract, the process is a labyrinth of documentation, regulation, and manual effort. Agentic AI can transform this by acting as a digital force multiplier for program managers and contracting officers.

Vision: Imagine a future where an AI agent, tasked with the objective of "procuring a new transportation system," automatically ingests and analyzes the Statement of Work. It could then search for and identify potential vendors, synthesize past contract performance data, and generate a preliminary draft of the Request for Proposal (RFP) for human review. It could also proactively alert the team to potential legal or regulatory concerns before they arise. Finally, it could analyze RFP responses to rapidly down-select partners whose propositions are closest to requirements. Once proposals are submitted, the same system could evaluate them at scale, assess risk, analyze schedule realism, and compare cost structures to historical benchmarks.

Impact: By automating the most tedious and time-intensive phases, Agentic AI can significantly reduce acquisition timelines and procurement costs. This would enable the DOW to acquire and field critical capabilities at a speed that is competitive with peer adversaries.

Application 2: Predictive Supply Chain Management

The defense supply chain is global, interconnected, and highly vulnerable. A disruption in one part of the world — a natural disaster, a geopolitical event, or a cyberattack — can have cascading effects that compromise operational readiness.

Vision: A set of AI agents could be deployed to continuously monitor the health of the supply chain. It could use real-time data from logistics networks, news feeds, and even weather patterns to predict potential disruptions before they occur. Upon identifying a potential risk, such as a port closure, an agent could proactively propose alternative shipping routes, identify secondary suppliers, and re-optimize logistics for key components, all while maintaining strict security protocols.

Impact: This proactive approach to supply chain resilience would ensure continuous readiness, minimize costly disruptions, and provide senior leaders with unprecedented transparency and foresight into their industrial base.

Application 3: Increased Quality Assurance Effectiveness

Agentic AI streamlines various quality assurance processes, including technical data package completeness review, first article inspection documentation, certificate of conformance validation, and discrepancy report processing. It performs comprehensive checks in minutes, identifies inconsistencies, verifies approvals, analyzes inspection reports, validates certificates, and automates discrepancy resolution — significantly reducing the time and effort required by engineers and quality assurance specialists.

Vision: By leveraging advanced AI capabilities to ensure accuracy, consistency, and efficiency in documentation and inspection processes, Quality practitioners will be able to focus on high-value activities. By learning from historical data and identifying subtle indicators of potential quality issues, Agentic AI can enhance the overall quality management system, making it more proactive and reliable.

Agentic AI for quality process extends beyond enhancing Quality Management Systems and be an integral part of the part design process. Throughout each step of the part design process, the Agentic AI system will continually analyze the part specifications, as well as geometric dimensioning and tolerancing, allowing each feature to continually provide guidance to design engineers for manufacturing viability, and the probability of meeting those requirements. This change will increase the first yield of First Article Inspections and reduce non-conformances at the design stage.

Impact: Implementing Agentic AI will reduce the time spent on manual reviews, minimize human errors, and improve the accuracy of quality checks. This will lead to faster issue identification and resolution, better compliance with contract requirements, and ultimately higher quality products. The system's ability to learn and adapt will ensure continuous improvement in quality assurance processes, reducing the cost of poor quality and providing long-term benefits to organizations.

Application 4: Accelerated Research & Development (R&D) and Innovation

The "valley of death" is a well-known phenomenon in defense R&D, where promising new technologies fail to transition from the laboratory to the field. One key reason is a lack of seamless collaboration between DOW labs, private-sector innovators, and end-users.

Vision: An AI agent, with the goal of "accelerating innovation in hypersonic technology," could act as a digital matchmaker. It could analyze a DOW lab's research findings, cross-reference them with academic publications and private-sector patent databases, and identify potential industry partners with complementary capabilities. It could even synthesize technical papers into clear, concise summaries for non-technical leadership, highlighting the most promising applications.

Impact: By bridging the gap between research and practical application, Agentic AI can dramatically increase the speed of innovation, ensuring the United States maintains its technological superiority.

Application 5: Dynamic Strategic Planning and Resource Allocation

In a dynamic global environment, strategic planning is no longer a static, once-a-year event. Senior leaders need the ability to rapidly model scenarios and make informed decisions about resource allocation in real-time.

Vision: An AI agent, given the objective of "optimizing the FY2026 budget for strategic readiness," could analyze vast datasets of historical spending, operational readiness reports, and geopolitical intelligence. It could generate a series of data-informed strategic plans, each with a detailed risk-reward analysis. A senior leader could then task the agent to run "what-if" scenarios, such as "What is the impact of a 10 percent reduction in the sustainment budget for a specific weapon system on its operational availability?" The agent would then provide a quantified, data-driven answer.

Impact: This capability would empower senior leaders to move beyond reactive budget cycles and engage in truly proactive, data-informed decision-making.

V. Navigating Risks and Ensuring Responsible Deployment

The integration of Agentic AI is not without risk. The move from human-in-the-loop (where the human must approve every action) to human-on-the-loop (where the human is an overseer and can intervene if needed) requires a new level of trust and accountability. Key risks to be addressed at a policy level include:

- Accountability: We must establish clear lines of responsibility. When an AI agent makes a decision that leads to an adverse outcome, it must be clear who is accountable — the programmer, the operator, or the organization.
- Explainability: AI systems can be like a black box. It is important for humans to understand the process through which the AI arrived at a specific decision.
- Ethics and Bias: All AI systems are susceptible to bias. Robust governance frameworks must be established to ensure that agents are free from bias and that their actions are aligned with human values and principles. Existing federal frameworks, including the AI Bill of Rights and M-25-21, "Accelerating Federal Use of AI through Innovation, Governance, and Public Trust", direct agencies to evaluate document migrations, and ensure alignment with established ethical norms. Agentic AI will require even more rigorous validation due to its ability to initiate or sequence actions autonomously.
- Security: Autonomous systems present a new attack surface for adversaries. The Office of Management and Budget (OMB) memorandum M-25-22, "Driving Efficient Acquisition of Artificial Intelligence in Government" underscores that systems procured by the federal government, or used by the federal government, must adhere to secure by design practices that undergo continuous monitoring. For defense applications, this may require treating agentic models as operational systems, not just software, requiring safeguarded architecture and procedural safeguards. Safeguarding against the unique cyber threats posed by these systems will be a top priority. These are not insurmountable challenges. They are policy and governance questions that can be solved through a methodical, phased approach. The key to mitigating these risks is a commitment to transparency, explainability, and the foundational principle that humans will retain the ultimate authority and oversight.

VI. Conclusion

The adoption of Agentic AI is a strategic imperative. It is about building a more adaptive, resilient, and strategically superior partnership between the DOW and the DIB. The future of defense decision-making is not a human-versus-machine proposition; it is a collaborative endeavor where Agentic AI acts as a force multiplier for human intellect, accelerating and improving decision-making in a rapidly evolving world.