

6 May 2026

Authors:

Stephen W. Roe

Dr. Nathan Turnipseed

Ryan Novak

Adam Bouffard

Wilson Miles

Adam Kitay

Kevin Forbes

Christopher Barlow

MITRE

ETi EMERGING
TECHNOLOGIES
INSTITUTE

NDIA



Keeping Humans in Command

Preserving Human Agency in AI-Driven
Federal Acquisition

ETi
NDIA

MITRE

SOLVING PROBLEMS
FOR A SAFER WORLD®

OVERVIEW



01

The AI-to-AI Paradox and Why Human Agency Matters

Exploring the paradox of AI systems making high-stakes decisions about each other—and why human agency remains essential.



02

Risks of AI-to-AI Acquisition

Examining key risks including bias amplification, opaque decision-making, compliance gaps, and loss of accountability.



03

The Human Agency Assurance Framework

Introducing the HAAF a framework to ensure human agency, transparency, and accountability in AI-to-AI acquisition.



04

How the HAAF can be Implemented

A practical roadmap for integrating the HAAF into acquisition processes, tools, and organizational workflows.



05

Real World Impact and Recommendations

Real world implications, measurable benefits, and clear recommendations for leaders and practitioners.

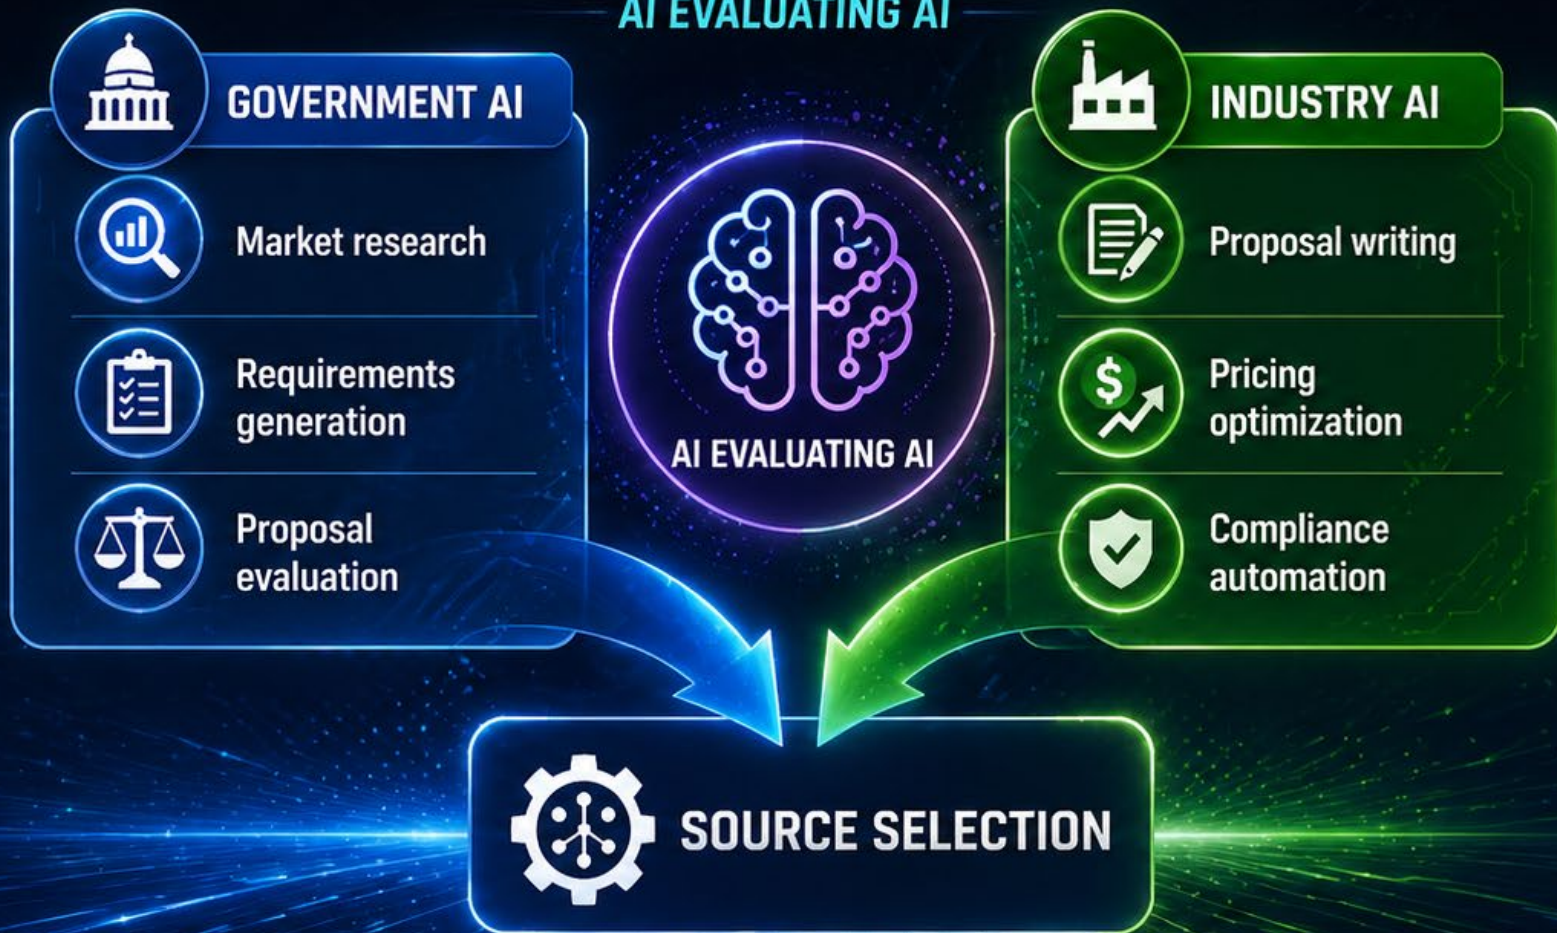


HUMAN AGENCY ENSURES AI SERVES OUR MISSION

The AI-to-AI Paradox and Why Human Agency Matters

TWO ECOSYSTEMS CONVERGING

AI EVALUATING AI



HUMAN AGENCY PROVIDES:



ETHICAL JUDGMENT



CONTEXTUAL UNDERSTANDING



ACCOUNTABILITY



MISSION ALIGNMENT



WARFIGHTER PERSPECTIVE

RISKS OF AI-TO-AI ACQUISITION



THESE RISKS INTERACT AND REINFORCE EACH OTHER,
COMPOUNDING INTO **SYSTEMIC RISK**.

HUMAN AGENCY ASSURANCE FRAMEWORK (HAAF)

ENSURING HUMAN AGENCY IN EVERY AI-DRIVEN DECISION

HAAF FOR SOURCE SELECTIONS

TIERED TOOL SETS



**INDUSTRY AI
PROPOSAL GENERATION**



**GOVERNMENT AI
EVALUATION TOOLS**



**HUMAN DECISION
AUTHORITY**



TIER	OVERSIGHT
 TIER 0	AUTOMATION PROHIBITED AI may support drafting or retrieval but cannot recommend or execute decisions.
 TIER 1	HUMAN-IN-THE-LOOP AI may generate recommendations, but a human must review the rationale and explicitly approve the output before it is used.
 TIER 2	HUMAN-ON-THE-LOOP AI may execute bounded tasks with continuous monitoring and rapid intervention capability.
 TIER 3	ADMINISTRATIVE AI AI may operate with minimal oversight where errors are easy to detect and correct and have negligible impact.



INCLUDES: Oversight thresholds, assurance gates, governance, policy and guidance, and workforce readiness.



THE RIGHT LEVEL OF HUMAN OVERSIGHT. AT EVERY LEVEL.

HOW HAAF WORKS

OPERATIONAL MODEL

DECISION FUNNEL



HUMAN APPROVAL POINTS



THREE PRACTICAL MECHANISMS

1



OVERSIGHT TIERS

- Tier 0: No automation
- Tier 1: Human-in-the-loop
- Tier 2: Human-on-the-loop
- Tier 3: Admin AI

2

ASSURANCE GATES



VALIDATE



MONITOR



AUDIT

3



GOVERNANCE CONTROLS

- AI oversight boards
- Contract requirements
- Explainability + traceability



HAAF ENSURES HUMAN JUDGMENT, ACCOUNTABILITY, AND TRUST AT EVERY DECISION POINT.

REAL WORLD IMPACTS

AI-DRIVEN. HUMAN-LED. MISSION-FOCUSED.



ACQAIRE

ACQUISITION AI RESEARCH

STRONGER DECISIONS.
FASTER INSIGHTS.
ZERO COMPROMISE ON

HUMAN AGENCY.

ETI
NDIA
MITRE
PARTNERING FOR
A SAFER FUTURE



AI-ENABLED MARKET RESEARCH AT SCALE:

The VA used a secure, bounded generative AI micro-tool to process 75+ extensive RFI responses, reducing analysis time from 8–12 weeks to 2–3 weeks while avoiding analyst fatigue and preserving rigor.



DECISION-SUPPORT:

The AI performed structured data extraction, thematic clustering, and summary generation under Tier 1 Human-in-the-Loop oversight, with all outputs treated as preliminary and subject to mandatory human validation.



HUMAN-LED VALIDATION AND INSIGHT:

Acquisition professionals conducted calibration, thematic analysis, and contextual interpretation identifying gaps and refining requirements (PWS, evaluation criteria) beyond what AI alone could infer.



TRACEABILITY, ACCOUNTABILITY, AND IMPACT:

All decisions were documented and auditable, ensuring defensibility and alignment with acquisition standards; AI functioned strictly as a force multiplier, enhancing speed and analytical depth while preserving human judgment and accountability.



FFRDC

FFRDC-LED AI EXPERIMENTATION ENVIRONMENT:

MITRE Corporation's ACQAIRE initiative provides a controlled environment to test acquisition focused AI under human-in-the-loop governance, demonstrating responsible integration across government, industry, and FFRDC contexts.



AI-AS-TOOLS FOR DECISION SUPPORT ONLY:

AI is deployed as narrowly scoped tools that surface information not conclusions, reinforcing human authority and mitigating automation bias.



EVIDENCE TAGGING AND TRACEABILITY:

AI outputs are explicitly linked to source material (text, page, criteria), enabling audibility, transparency, and protest defensibility, while reducing risks of hallucination or misinterpretation.



COGNITIVE AUGMENTATION WITH PRESERVED JUDGMENT:

AI reduces evaluator workload (e.g., highlighting inconsistencies, compliance gaps) to accelerate analysis, but humans retain full responsibility for interpretation, tradeoffs, and final decisions, aligning with HAAF principles of accountability and oversight.



DELIVERED FOR
VETERANS.
POWERED BY
HUMAN AGENCY.



ACQAIRE

MITRE'S ACQUISITION
AI RESEARCH
ENVIRONMENT.



STRENGTHENING HAAF OVERSIGHT

ACTIONABLE RECOMMENDATIONS FOR TRUSTED AI IN ACQUISITION

01



CIRCUIT BREAKER FOR REAL-TIME AI FAILURES

Establish predefined triggers (e.g., low confidence scores, drift detection, protest signals) that automatically pause AI-assisted acquisition processes and revert to manual evaluation, closing a critical gap in HAAF for in-process failure response.



DETECT. PAUSE. PROTECT.

02



INDEPENDENT AI ACQUISITION AUDITOR

Create a neutral oversight function (e.g., within Government Accountability Office) with authority to audit AI use in acquisition, analyze protest trends, and publish findings—ensuring external accountability beyond internal governance structures.



INDEPENDENT. OBJECTIVE. ACCOUNTABLE.

03



POLICY INTEGRATION INTO DOW ACQ FRAMEWORKS

Embed HAAF oversight tiers into DoW Instruction 5000.02 to formally govern AI usage across MDAPs, OTAs, and rapid pathways, aligning AI oversight with existing acquisition lifecycle controls.



INTEGRATED. ALIGNED. GOVERNED.

04



MANDATORY SUNSET AND REVALIDATION

Require periodic revalidation of AI tools with automatic suspension if not completed, addressing model drift and obsolescence to ensure long-term reliability, defensibility, and trust in evolving acquisition environments.



REVALIDATE. RELY. REPEAT.



STRONGER OVERSIGHT. SMARTER DECISIONS. MISSION SUCCESS.

MITRE

ETI

NDIA

Keeping Humans in Command

Preserving Human Agency in AI-Driven Federal Acquisition

NDIA Emerging Technologies Institute
The MITRE Corporation
May 2026

QUESTIONS



MITRE continues to work with organizations such as **NDIA/ETI** and our sponsors to identify how to leverage AI in acquisition across the ecosystem with projects such as **ACQAIRE**.

« PLEASE REACH OUT TO THE FOLLOWING IF YOU WOULD LIKE TO DISCUSS OR HAVE ANY FURTHER QUESTIONS. »»



Adam Bouffard

✉ abouffard@mitre.org



Ryan Novak

✉ rnovak@mitre.org



Wilson Miles

✉ wmiles@ndia.org



“Ultimately, the objective is not to slow the adoption of AI within federal acquisition, but to **guide it**. The warfighter’s need for speed must be balanced with institutional safeguards that confirm acquisition decisions remain defensible, ethical, and strategically sound. When implemented responsibly, AI can serve as a **force multiplier for human expertise** rather than a substitute for it. **Keeping humans in command** confirms that technological acceleration **strengthens**, rather than undermines, the trust and accountability upon which effective defense acquisition depends.”



MITRE

ETI

EMERGING
TECHNOLOGIES
INSTITUTE

NDIA



STRONGER OVERSIGHT. SMARTER DECISIONS. MISSION SUCCESS. »»»