

OCTANT **ADVISORY**

AI REALITY SERIES · MAY 2026

WHITE PAPER

AI Reality for GovCon and Defense Contractors

How federal customers are actually evaluating AI maturity, and why most primes are not yet prepared.

DR. IAN MCCULLOH

Co-Founder & Chief AI Strategy
Officer

MARIA

CHALOUX

Co-Founder & Managing
Partner

OCTANTADVISORY.COM

EXECUTIVE SUMMARY

Compliance produces documentation. Customers are buying capability.

AI maturity is quietly becoming a factor in which firms win and lose federal contracts, and the way it is being judged catches even sophisticated, fully-compliant primes off guard.

Federal customers are not waiting for a unified standard. They are forming real-time judgments (during source selection, oral presentations, past-performance review, and contract execution) about whether a contractor can be trusted to deploy AI responsibly at scale.

Most firms are investing to be compliant. That investment is necessary, but it is not what is being evaluated. A prime can align to the NIST AI Risk Management Framework, hold CMMC 2.0 certification, and adopt DoD Responsible AI principles, and still fail to demonstrate AI maturity in the way a federal customer experiences it. Maturity is not read off a compliance artifact. It is inferred from how the organization actually operates: how decisions get made, where authority sits, how risk is handled in real time, and whether outcomes are measurable and defensible.

Most primes are optimizing for compliance artifacts. Federal customers are selecting for organizational capability.

That divide is beginning to surface in source-selection scoring, technical evaluations, and CPARS, and the firms on the wrong side of it tend to lose AI-implicated work without a clear reason ever appearing in the debrief.

This paper introduces the Octant AIR Index™ (the AI Reality Index): a practical framework for evaluating enterprise AI maturity across three lenses: AIR-O (organizational), AIR-T (technical), and AIR-L (leadership). Each is scored on a five-level maturity scale. It focuses on AIR-O, the organizational conditions that determine whether AI investment converts into contract-winning performance, and it gives leaders a self-diagnostic they can run against their own firm today.

01

The evaluation environment is changing faster than the rules

A Chief AI Officer at a prime today is expected to demonstrate AI maturity against a fragmented and fast-moving set of expectations:

- Federal customer evaluation criteria: source selection, oral presentations, performance reviews, and CPARS
- NIST AI Risk Management Framework and ISO/IEC 42001

- DoD Responsible AI Strategy and DoD Directive 3000.09 on autonomy in weapon systems
- CMMC 2.0 and its cybersecurity assurance regime
- ITAR and EAR export controls on AI models, training infrastructure, and technical data
- DFARS data rights frameworks governing AI-relevant data and IP
- Emerging AI-specific contract clauses entering FAR- and DFARS-based agreements

The ground underneath these requirements is still moving. Executive Order 14110 (October 2023) was rescinded and replaced by Executive Order 14179 in January 2025, reorienting federal AI policy toward accelerating adoption and removing barriers. The July 2025 AI Action Plan, "Winning the Race," set that direction across agencies. The FY2026 NDAA's Section 1513 signaled a CMMC-style assurance regime for AI, and CMMC 2.0 itself has moved from final rule into its phased rollout. The throughline is unmistakable: requirements are proliferating and shifting, not converging.

These frameworks were never designed to operate as a single system. They sit at different levels (principles, compliance requirements, contractual obligations, and operational expectations), so the documentation surface is large, the guidance is evolving, and interpretation is inconsistent across customers and program offices.

Federal customers resolve that ambiguity in practice. They do not select the most compliant contractor. They select the contractor they believe can be trusted to operate AI under uncertainty.

02

The core failure mode: compliance does not produce capability

A prime can demonstrate alignment across every major framework and still fail to demonstrate AI maturity. This is not theoretical; it shows up in source-selection scoring, technical demonstrations, and past-performance review.

The pattern is consistent. Compliance produces documentation. Customers are evaluating capability. A firm may have NIST-aligned policies, ISO-referenced governance structures, and CMMC-certified environments, yet still lack enforceable decision rights, leadership authority to intervene, and operational accountability for AI outcomes.

Documentation demonstrates intent. Organizational structure determines outcome.

Federal customers have learned to read past the artifacts to the capability underneath. The governing principle is blunt: if governance authority cannot be exercised, it does not exist. That is why firms that appear fully compliant on paper still fail in competitive environments.

03

Where AI maturity actually shows up

In the absence of a unified standard, customers fall back on what they can observe: how the organization behaves under pressure. That observation is not abstract. It happens at specific, identifiable moments in the acquisition lifecycle.

EVALUATION MOMENT	WHAT THE CUSTOMER OBSERVES	WHAT THEY INFER
Source selection	How AI governance is described and who is named as holding decision rights	Whether authority is real or narrative
Oral presentations	Whether named personnel can answer governance questions unscripted	Depth versus theater
Technical demonstrations	Whether AI behavior is monitored, explainable, and controllable	Technical maturity under real conditions
Past performance / CPARS	How the firm handled AI-related incidents, drift, or change	Operational accountability over time
Contract execution	Whether governance authority is actually exercised when systems behave unexpectedly	Trustworthiness under uncertainty

Consider an illustrative composite. A top-tier prime arrives at an oral presentation for an AI-implicated re-compete with full NIST AI RMF alignment and current CMMC certification. An evaluator asks the named AI lead a simple question: who has the authority to halt an AI system in production, and has that authority ever been exercised? The answer reveals that the authority formally sits with a governance committee that has never convened on an operational matter. The documentation is impeccable; the organizational capability is not. The score reflects the capability, not the documentation.

04

The Octant AIR Index™

The Octant AIR Index™ (the AI Reality Index) reads enterprise AI across three integrated lenses:

- **AIR-O (Organizational):** enterprise design, governance, and operating model
- **AIR-T (Technical):** data, infrastructure, security, and deployment capability
- **AIR-L (Leadership):** the individual leadership capability to drive transformation

Each lens is scored on a five-level maturity scale, so a firm's maturity can be described precisely rather than asserted:

LEVEL	STAGE	WHAT IT MEANS
Level 1	Ad hoc	AI activity is isolated, undocumented, and personality-dependent
Level 2	Emerging	Policies and pilots exist but authority and outcomes are unclear
Level 3	Established	Governance is defined and repeatable across the enterprise
Level 4	Demonstrated	Capability is exercised, auditable, and visible to customers
Level 5	Optimized	AI maturity is continuously measured, improved, and tied to performance

Most organizations focus on AIR-T. Federal customers are increasingly weighting AIR-O. Most firms have AI capability; few have the organizational conditions required to deploy it at scale. The remainder of this paper focuses on AIR-O.

05

AIR-O: the eight conditions, and how to read your own firm

AIR-O defines eight structural conditions required for AI to deliver outcomes. They are diagnostic, not aspirational. The table below is a self-check: for each condition, weak and strong look distinctly different to an evaluator, and to a customer watching you operate.

AIR-O CONDITION	WHAT WEAK LOOKS LIKE	WHAT STRONG LOOKS LIKE
01 · Strategic Mandate	AI appears in strategy decks but not in resourced plans or named outcomes	AI ambition is board-defined, tied to specific mission or contract outcomes, and funded
02 · Executive Alignment	Leaders describe AI priorities differently when asked separately	The C-suite articulates a consistent AI thesis, including its trade-offs
03 · Governance & Risk Authority	Decision rights exist in a policy document, but no one has ever halted a system	Named owners hold enforceable authority that has actually been exercised
04 · Technology Strategy & Governance	AI tools are adopted team-by-team without architectural or risk integration	AI workloads are inventoried, tiered by risk, and governed against architecture and compliance
05 · Operating Model & Structure	AI sits in a center of excellence with no line authority	

AIR-O CONDITION	WHAT WEAK LOOKS LIKE	WHAT STRONG LOOKS LIKE
		AI capability supports both proposals and delivery, with clear interfaces and ownership
06 · Leadership Capability	The CAIO has a title but no budget, staff, or authority to intervene	AI leadership has authority commensurate with accountability, and bench depth behind it
07 · Workforce & Culture	Training is one-off and adoption is uneven	The workforce has role-based readiness and a feedback loop back into governance
08 · Value Measurement & Learning	AI benefits are asserted, not baselined	Outcomes are measured against pre-deployment baselines and tied to performance

06

What good looks like: the Level 4 prime

On the AIR-O five-level scale, a Level 4 prime (Demonstrated) does not describe its AI capability. It demonstrates it through repeatable, auditable behavior across all eight conditions:

- AI is embedded in enterprise strategy with defined, resourced outcomes
- Leadership is aligned on priorities and trade-offs
- Governance authority is enforceable, and has been exercised
- Risk and technology architecture are integrated, not parallel
- The operating model scales from proposal to delivery
- Leadership has both authority and depth
- The workforce is ready and its experience feeds governance
- Outcomes are measured against baselines and tied to performance

This capability is visible in proposal environments, technical demonstrations, contract execution, and CPARS. It is not a narrative. It is evidence, and it is exactly what federal customers are now looking for.

07

Implications for GovCon leadership

For the Chief AI Officer. If you cannot influence decisions across the enterprise (budget, architecture, and the authority to intervene), AI maturity cannot be demonstrated, no matter how complete the documentation. Read each AIR-O condition as a question about your own standing.

For the Chief Risk Officer. Governance authority that is never exercised is indistinguishable, to a customer, from governance that does not exist. The most effective move is to make decision rights real and to be able to show where they have been used.

For Business Unit and Capture Leaders. If AI capability cannot be demonstrated in a proposal or an oral presentation, it will not influence the award. The eight conditions are a checklist for what evaluators are actually reading.

For the General Counsel. If data rights, export controls, and AI governance are not structurally integrated, the exposure will surface during execution, the most expensive place to discover it.

For executive committees and boards. The eight-condition structure provides a shared vocabulary for governing AI readiness as an enterprise capability rather than a technology project.

08

Conclusion

AI maturity is no longer a future capability. It is a present selection criterion.

Federal customers are not evaluating whether you meet standards. They are evaluating whether your organization can be trusted to operate AI under real conditions. That distinction will increasingly determine which firms win AI-implicated work, and which do not.

RUN THE DIAGNOSTIC

An AIR-O snapshot of your firm

Octant offers a structured, confidential read of your firm against the eight organizational conditions, benchmarked to how federal customers evaluate AI maturity. To request one, contact the Octant team at octantadvisory.com.

ABOUT OCTANT ADVISORY

Octant Advisory is an AI transformation advisory and executive search firm focused on the structural conditions that determine whether enterprise AI investments deliver outcomes. The firm was founded in 2026 by Dr. Ian McCulloh and Maria Chaloux. Octant's framework, the Octant AIR Index™ (the AI Reality Index), reads organizational, technical, and leadership maturity across a single integrated model. For inquiries about the AIR Index, the AI Reality diagnostic, or executive advisory engagements, visit octantadvisory.com.

This paper is provided for general informational purposes and reflects Octant Advisory's framework and professional judgment as of May 2026. It does not constitute legal, regulatory, or compliance advice. Regulatory requirements referenced here are evolving; readers should consult qualified counsel before making decisions based on them. © 2026 Octant Advisory · octantadvisory.com