G1VE TECHNOLOGIES

Federated AI Without Governance Is Going to Cost Us The War We Are Building For

W.B. Carrington, MBA, LSSMBB, CTPRA, GRCA, GRCP

Sep 9, 2026

Original source

https://www.linkedin.com/pulse/federated-ai-without-governance-going-cost-us-war-we-w-b--p3nee/

I. The Cost of a Bridge With No Deck

Federated AI without sovereign governance is going to cost us the war we are building for. That thesis is not alarmism; it is the September 2026 forecast. The United States Department of War has now published the most recent iteration of its AI use-case inventory, the Office of Management and Budget of management has refreshed the consolidated federal AI inventory through the M-25-21 reporting cycle, and the total of agency-reported AI use cases has cleared the 2,200-notional mark across the federal enterprise. The United States Department of War AI Inventory is approaching 1,000 unclassified use cases on its own, with the heaviest concentration logged in administrative and logistics processes — not the warfighting portfolios the FY26 National Defense Strategy pre-supposes it should be serving. We have built the bridge. We have not poured the deck. We are about to drive the first brigade across it, live, on the operational network.

Three follow-on facts:

1. National Institute of Standards and Technology (NIST) AI 600-1 — the Generative AI Profile companion to the AI Risk Management Framework — is in final publication cadence per the live National Institute of Standards and Technology (NIST) landing page .
2. The GSA AI Center of Excellence and the GovAI Coalition have begun publishing model-card templates that are about as tightly-spec'd as a real model card can be, which means the inventory paper product is becoming queryable, auditable, and increasingly contractually significant.
3. The FY27 NDAA policy draft is treating federated AI as a default operating mode for combatant commands and for the IC. That is the right call. It is also ungoverned.

...We have about eighteen months to fix it before the bridges we are pouring today become the kill chains we cannot afford tomorrow.

II. The State Of The AI Inventory — Quantity Is Up, Sovereignty Is Down

Inventory growth is the easy half of the problem. The first unclassified United States Department of War AI Inventory report was tacked to the FY19 NDAA Section 1051 reporting requirement; the inventory has grown in scope and depth with each FY cycle. The current OMB M-25-21 inventory portal — successor to M-24-10 — is the consolidating instrument for federal agency use cases (OMB — M-25-21). The Brennan Center wait-time analysis on the December 2025 inventory cycle gives a sober readout:

- The bulk of reported use cases remain administrative, with a large fraction in operational and logistics support — and the largest aggregate categories are still in the back-office.

That is not a flaw in the United States Department of War AI team. It is a flaw in how the United States Department of War measures AI work product when operational portfolios behave carefully about disclosing their AI-funded capability lines.

Sovereignty is the harder half. When a federally-funded AI model depends on — or is hosted by — a covered foreign-country training pipeline, or a covered foreign-country inference compute stack, the sovereignty margin is negative on day one. The FY24 NDAA Section 142 reporting requirements are now being flowed through the FY25 and FY26 follow-ons, and the CDAO has standing authority to require additional inventory metadata from program offices.

Combined with the U.S. Department of the Treasury Outbound Investment Security Program final rule (28 Oct 2024, effective 2 Jan 2025) and the FY26 NDAA's heightened ownership and control disclosure provisions, every model now has a fiduciary owner of record. If your office cannot stamp that line, your model is not in the inventory — it is a liability.

III. The Governance Stack in 2026 — Everything You Need, Nothing Wired Together

The United States Department of War AI Ethical Principles are five years old. They still hold. The Responsible AI (RAI) Strategy and Implementation Pathway issued in 2024 laid the procedural framework. The DoW Super Intelligence Office - SIO (formerly CDAO) 's RAI Toolkit is the operational instrument that PMs, COCOM J6s, and OSBPs are now expected to apply at acquisition drafting, prototype review, and model card publication. The OMB M-25-22 policy on federal AI risk management (successor to M-24-18) sits across all instruments and provides the cross-government overlay. The National Institute of Standards and Technology (NIST) AI RMF 1.0 is the baseline; AI 600-1 is the GenAI extension. The GSA AI Center of Excellence and the GovAI Coalition standardized the model card and federal AI procurement language. CAIOs are now named across every covered agency.

The sovereign question is: who owns the wiring? In practice, the wiring is owned by every office that touches the inventory product and no office that owns its lifecycle. That is the gap the FY26 NDAA has started to address by formalizing federated model governance, by directing inventory model-card publication to be auditable, and by setting thresholds for sovereign compute, sovereign data, and sovereign inference.

If your office is not yet running an end-to-end RAI review on every model that touches operational data — specifically the foreign-exposed training classes that come through the DoW Super Intelligence Office - SIO (formerly CDAO) 's authorized cloud tenants — you are not compliant with the policy stack. Your §142 report just shifted from "approximately accurate" to "non-existent."

IV. Federated By Design — The Operational Logic The FY26 Fight Demands

Federated AI is the operational logic the FY26 fight demands. The United States Department of War cannot afford to centralize every model on a single classified tenant; the bandwidth, the classification tiering, and the COCOM-to-COCOM operating tempo make a central stack operationally untenable for the forces that need it. The NATO Science & Technology Organization (STO) is an early example of the federated doctrine in production. The Defense Innovation Unit (DIU) 's DAILi-ITA AI Cluster work, the U.S. Army DEVCOM Army Research Laboratory 's GARD program working in partnership with U.S. Naval Research Laboratory and

the Air Force Research Laboratory , and Project Maven's federated model-sharing exercises all demonstrate the operational logic.

The NOAA: National Oceanic & Atmospheric Administration and the Centers for Disease Control and Prevention both maintain federated AI architectures for hurricane forecasting and communicable disease surveillance, respectively, and the academic base cases for federated model training and inference over sovereignty-respecting partitions come from foundational work published in arXiv and SPIE going back to 2019. The script of what works is consistent:

- Distributed data stays distributed,
- Sovereign inference happens at the source, and,
- The federated model shares gradients only.

...That script requires three things the current governance stack does not yet guarantee:

1. Model-card traceability across federation boundaries,
2. Auditable lineage on every gradient update, and,
3. A kill-switch that can isolate a sovereign inference authority when contamination is detected.

Three live United States Department of War programs illustrate the gap:

1. Project Maven's expanding federated model sharing at Anduril Industries × Palantir Technologies ' Lattice integration,
2. NASA - National Aeronautics and Space Administration 's / the NOAA: National Oceanic & Atmospheric Administration 's hurricane models, and,
3. The U.S. Army DEVCOM Army Research Laboratory 's GARD federated output-protection research line.

Each demonstrates the operational logic. None yet ships a complete governance stack. That is the asymmetric vulnerability we are running into a Great-Power pacing threat.

V. Zero Trust AI — The Overlay That Finally Meets the Stack

Zero Trust is no longer the CISO's slide deck; it is the United States Department of War's operating posture. The United States Department of War's Zero Trust Strategy and Roadmap issued in 2022 was an opening shot; DTM 25-003, published in 2025, conferred the strategy onto IT and OT for both cloud computing and on-prem environments, with effective date obligations running from late 2025 onward.

The strategy has five lines of effort; AI workloads now sit explicitly on Lines 1, 3, and 5 workload identity, and on Lines 2 and 4 for cross-domain governance. The National Institute of Standards and Technology (NIST) 's SP 800-53 Rev 5 and the Office of Management and Budget 's M-22-15 anchor the federal overlay; the FY27 ZT Strategy target is full zero trust at the FY27 program-of-record horizon. The PfMO and the Chief ZT Officer, Randy Resnick , at the United States Department of War are now the execution layer.

AI workloads are where Zero Trust meets the stack at last. The DoW Super Intelligence Office - SIO (formerly CDAO) 's AI T&E and the AI T&E Lab are doing the heavy lifting on model-card

traceability, audit trails, and tamper-evident training, and these tools are now being flowed into the Zero Trust execution discipline (NIST SP 800-53 control overlays). The U.S. Department of Veterans Affairs branch has published the U.S. Department of Veterans Affairs ' Federated Mesh reference architecture as a use case for federated Zero Trust AI at agency-transcending scale.

The DoW Super Intelligence Office - SIO (formerly CDAO) 's AI Test & Evaluation platform is now the operational ground where model quality, sovereign-data exposure, and adversary-data-poisoning rubrics meet the OT test bench. The next eighteen months will tell whether the operational AI workloads the FY26 NDS demands can transit ZT5 with auditable lineage — or whether the legacy 2018 IT posture reasserts itself. If your program office is not running ZT-aligned AI workload identity today, your program is not in the FY27 NDAA target postrure.

VI. Procurement & PoRs — Where The Money Actually Flows

The money has already moved. JWCC — Joint Warfighting Cloud Capability — was awarded to Amazon Web Services (AWS) , Google , Microsoft , and Oracle in December 2022, and total task-order obligating through Q3 FY25 sat at the $3.9 billion mark. The Defense Information Systems Agency August 2025 SAM.gov solicitation preview for the planned next-generation enterprise cloud program is already signaling what comes next. Maven Smart System (MSS) — the United States Department of War's primary AI-portfolio acquisition vehicle — has been on a steady awards trajectory.

The most recent MSS contract modification — the $795 million modification to Anduril Industries ' × Palantir Technologies ' existing contract vehicle under W911QX-24-D-0012 — was logged 21 May 2025; the US Army awarded $480 million in May 2024 for prototype expansion; the US Army 's $99.8 million ancillary task order came in September 2024; the United States Marine Corps followed up with a $192 million enterprise-wide MSS-equivalent task order in September 2025. The Anduril Industries ' × Palantir Technologies ' Maven Lattice made its first joint commitment in December 2024.

Adjacent to MSS, the GSA 's AutoAuthority and the GSAi Pathfinder contracts are shaping the AI acquisition landscape, by focusing on automated AI procurement and AI-pathfinder commercial acquisitions that bypass the FAR Part 15 timeline. The FY26 NDAA also opens small-business AI ladders through the OSBP small-business set-aside pilot.

...The structural risk:

- The Capability Acquisition Executive (CAE) and Portfolio Acquisition Executive (PAE) models under the Acquisition Transformation Strategy are veering authority away from the PEO chair, and MSS-equivalent vehicle owners sit on those new chairs.

If your program office has not yet tied its AI acquisition lane into an MSS-equivalent vehicle, the FY27 funding cycle will not wait for you. The US Government Accountability Office has continued to flag persistent challenges in United States Department of War AI work-product procurement GAO-25-107933, and the FY27 NDAA United States Senate committee work is treating those gaps as blockers to the AI-close-deadline posture.

VII. The Three Bets for CDAOs and CMOs — Execute in the Next 90 Days

Three bets. Execute them in the next 90 days:

1. Bet One — Governed Artifact by FY26 close. Every CDAO-listed AI tool, every GSA Pathfinder intake, and every JWCC tenant AI tool must publish a model card and a 50-state lineage audit by FY26 close. This is the bare minimum for §142 reporting. The model card should be auditable; the lineage audit should pass run-of-paper. The DoW Super Intelligence Office - SIO (formerly CDAO) RAI Toolkit is the publication vehicle. Anything less than full model-card publication is a §142 defect waiting to crater the program on the FY27 review cycle.
2. Bet Two — Operationally Mapped Inventory. Every DoW Super Intelligence Office - SIO (formerly CDAO) -listed tool must be operationally mapped to the COCOM war-room it serves. That mapping is what flows into the FY26 NDAA federation model, into the inventory analytics the DoW Super Intelligence Office - SIO (formerly CDAO) is required to publish under §142, and into the United States Department of War's Mission Assurance Construct's Day-1 inputs we covered in the last Securing Our Future editions. No operational mapping — no inventory integrity. No inventory integrity — no material line in the §142 report. That is how the US Government Accountability Office will find your program. You can bet on it.
3. Bet Three — Federated Acquisition Playbook. Adopt a federated acquisition playbook — distributed training, distributed inference, shared gradients only — across every new AI award that touches operational data. The Acquisition Transformation Strategy memos, the GSA 's AutoAuthority, and the JWCC tenant architecture are the operating vehicles. Pull this into the Acquisition Transformation memo's portfolio review now; it will become a mandatory operating discipline before FY27 close and a GAO test case before FY28. Federate responsibly. Govern aggressively. Ship the model.

✨ Call to Action: Step Into the Arena ✨

Most professionals will read about the future. A select few will shape it.
If you are serious about staying ahead of the curve in Artificial Intelligence, Regulatory Compliance, and Enterprise Risk, this is where you need to be.
📩 Subscribe to Future Trends & Insights: AI to gain direct access to:

1. Strategic foresight on AI Governance & Ethical Deployment
2. Real-world applications across National Security & Mission Assurance
3. Advanced GRC, TPRM, and Regulatory Compliance frameworks
4. Emerging Fintech and Risk Intelligence innovations

⚔️ The landscape is shifting—fast. The gap between those who understand and those who execute is widening.
This is built for professionals responsible for decisions that carry consequence—not commentary.
📅 Ready to operationalize these insights and move beyond theory?

☞ https://calendly.com/g1ve-technologies
Let us translate insight into measurable, defensible outcomes—and build capability, not just conversation.

🔍 About the Author

W.B. Carrington, MBA, LSSMBB, CTPRA, GRCA, GRCP is a National Security-aligned Artificial Intelligence Strategist/Tactician and Governance, Risk, and Compliance (GRC) Authority/Subject Matter Expert with over a decade of experience operating at the intersection of National Security, National Defense, Critical Infrastructure, Cybersecurity, AI Governance, Mission Assurance, Mission-Critical Systems, and Predictive Funding/Grant Intelligence.

As Managing Partner at G1VE Technologies, LLC and a Strategic Board Advisor, W.B. Carrington, MBA, LSSMBB, CTPRA, GRCA, GRCP architects and operationalizes enterprise-scale risk, compliance, and AI governance ecosystems across both Public and Private Sector environments—including Fortune 500/1000 organizations and Federal-aligned missions, where failure is not an option.

His work is centered on delivering Decision Advantage, strengthening Operational Resilience, and enabling organizations to navigate high-stakes, high-uncertainty environments with precision and confidence.

📡 Connect with W.B. Carrington, MBA, LSSMBB, CTPRA, GRCA, GRCP here to stay at the forefront of:

- Artificial Intelligence
- National Security
- GRC
- Cyber Risk
- Regulatory Strategy
- Mission Assurance