AI in Defense Manufacturing

A Management Framework for Governed AI Deployment Under CMMC, DFARS, and Export Controls

David W. Koran

CyberAB Registered Practitioner Advanced

ISO/IEC 27001 Auditor and Implementer

August 2026

Summary

Defense manufacturers face the same artificial intelligence pressure as every other business. Employees adopt AI tools because they help, software vendors add AI features to products already on the shop's computers, and primes, customers, and insurance carriers have begun asking AI questions in writing. One thing separates a defense shop from every other adopter: the business runs on regulated information. Federal Contract Information, Controlled Unclassified Information, and export-controlled technical data move through the same systems where AI tools want to work, which means every AI deployment decision in a defense manufacturer is a data routing decision, whether or not anyone treats it as one.

This paper is a management plan for making those decisions deliberately. Its structure is a staged deployment ordered by data class: Stage One deploys the uses that never touch protected information, promptly and under basic controls, because that is where the business banks value and builds the habit of governed use. Stage Two deploys AI into the business systems, quoting, planning, meetings, and customer communication, under a defined tool review. Stage Three addresses the uses that would touch CUI, and it begins from an honest premise: most commercial AI services cannot make the commitments that safeguarding requires today, so Stage Three is an analysis with three legitimate outcomes: a specialized environment, a locally operated capability inside the existing boundary, or a documented decision not to deploy. Export-controlled technical data receives its own treatment, because sending it to a cloud AI service raises questions that belong with export counsel before any tool is turned on.

Each stage leaves records: the data map, the approved tool list, the review files, and the decision records. The records are the plan's product as much as the deployments are, because they are what keep the company's existing representations true: the SSP, the SPRS score, and the answers already given to primes and carriers, while AI enters the operation. The plan is a companion to How to Build an AI Governance Program, which describes the full governance program for organizations of every kind; this paper is that program arranged for the realities of a defense manufacturing business.

1. Why This Is Different in a Defense Shop

Four classes of information organize this plan because they carry regulatory consequence in a defense manufacturer, and an AI tool that is harmless against one can be a violation against another. Public information, the website, marketing material, and published capability statements, carries no restriction. Federal Contract Information, information provided by or generated for the government under contract and not

intended for public release, carries the basic safeguarding requirements every federal contractor accepts. Controlled Unclassified Information, the drawings, specifications, and technical data marked and handled under the CUI program, carries the safeguarding requirements of NIST SP 800-171 implemented under DFARS 252.204-7012, with incident reporting obligations attached. Export-controlled technical data under the ITAR or the EAR carries restrictions on who may access it and where it may go that operate independently of everything above. These four are not the only information requiring protection: proprietary business and pricing information, personal and personnel information, privileged communications, and information restricted by customer agreements also require handling rules, and the AI use policy in Section 9 names them. The same prompt, typed into the same AI assistant, is routine against the first class and a reportable problem against the third or fourth, and the tool is identical in both cases; the difference lies entirely in the class of data involved.

The second difference is that a defense shop has already made promises. The company attests to safeguarding covered defense information. It maintains a System Security Plan describing an environment with a defined boundary. It has entered an assessment score in SPRS. It has answered prime contractor questionnaires and flowdown certifications, and increasingly it has answered AI questions on its cyber insurance application. Those representations describe the environment as it stood when they were made. An unreviewed AI service that begins processing, storing, transmitting, summarizing, or indexing covered information changes the environment the representations describe, quietly, and the company discovers the change at the worst possible time, in an assessment, an incident investigation, or a claim. A dated note belongs here: as of August 2026, CMMC implementation stands in Phase 1 following the July 2026 suspension of Phase 2, and the suspension did not suspend FAR 52.204-21, DFARS 252.204-7012, SPRS assessment, or the underlying safeguarding obligations. An AI service that processes FCI, CUI, or security protection data can change the company's CMMC scope, evidence, and affirmation basis under any version of the program, which is part of why this plan's records matter regardless of where the program review lands.

Stated plainly: the question in a defense manufacturer is never whether AI is useful. It manifestly is, in quoting, planning, drafting, and analysis, and the shops that deploy it well will hold a real advantage. The question is which information each AI tool can reach, and nobody can answer that from a product page. The remainder of this paper is a plan for answering it in order, banking the value that is safely available now, reviewing what needs review, and documenting the decisions the harder cases require.

2. The Kinds of AI, by Where the Data Goes

Artificial intelligence is sold to executives as one thing, and for the purposes of this plan it is five things, distinguished not by how the technology works but by where the data goes when the tool runs. That is the only technical distinction a defense manufacturer's management has to hold, because the delivery model determines the data path, and the data path is what the stages govern. The five models below cover nearly everything a manufacturer will encounter, and the first executive question about any proposed AI use is which of the five it is.

Delivery model	**What it looks like**	**Where the data goes**	**Where it fits in this plan**
Public consumer AI services	The free and personal-account tiers of the public generative AI chatbots	To the provider under consumer terms, frequently retained and usable for training	Stage One material only, under organizational accounts; controlled information never, which is the first-day rule
Commercial cloud AI on business tiers	The same class of services under organizational accounts on enterprise terms	To the provider under the tier's negotiated terms; commitments differ sharply by tier	Stages One and Two, subject to the tool review
Embedded AI features	AI arriving inside software already owned: the office suite, the ERP, the meeting platform	Along the host product's vendor chain: its model providers and subprocessors	Staged by the data in the host system; reviewed on arrival, disabled by default pending review
Government-foc used cloud environments	Cloud offerings built to hold covered information, with AI capabilities operating inside them	Within an environment designed for the data, under flowdown-compatibl e terms	Stage Three candidates, evaluated under the cloud service provider analysis

Delivery model	What it looks like	Where the data goes	Where it fits in this plan
Locally operated models	Open-weight models run on the shop's own hardware through local runtimes	Nowhere; the data stays inside the shop's environment	Stage Three candidates; the shop owns the capability, the upkeep, and the logging

The taxonomy matters because the same capability at two delivery models is two different risk propositions. A drafting assistant on a consumer account and the identical assistant on an enterprise tier differ in everything the review cares about: retention, training use, and contractual recourse, while differing in nothing the user sees. The confusion runs in both directions: executives overestimate the risk of local models, which keep data in the building, and underestimate the risk of embedded features, which arrive silently inside trusted software. The stages that follow assume this vocabulary, and Stage Three in particular turns on the difference between the last two rows of the table.

3. The Plan in One View

The plan is a staged deployment ordered by data class, and the two axes described in this section divide the work between them: the data a tool can reach establishes its stage, assigned by the most sensitive information involved rather than by how useful the tool sounds, and the operational consequence of its output sets the review depth, the approval authority, and the human oversight the deployment requires. Before Stage One, management establishes the mandate described in the companion paper as Step Zero: an executive sponsor, a named owner with authority, the scope, and a reporting rhythm. In a 50-person shop, this is one meeting and one page, and skipping it is the most common cause of plans that stall at the first disagreement.

Stage	Data classes	Example uses	Review and records
Stage One	Public information and approved low-sensitivity internal information; no FCI, CUI, export-controlled technical data, personal, privileged, or contract-restricted information	Marketing and website drafting, administrative writing, scheduling, general research under organizational accounts	Basic controls: approved accounts and tiers, the data rules, training. Records: approved list entries, acknowledgments
Stage Two	Two-A: business information with no FCI. Two-B: systems that process, store, or transmit FCI	Quoting and ERP features, meeting assistants for designated internal meetings, customer service drafting, document search over unrestricted repositories	Tool review per Section 6; for Two-B, FAR 52.204-21 applicability and scope documentation in addition. Records: review file and approval decision record per tool and embedded feature
Stage Three	CUI and export-controlled technical data	Engineering document analysis, technical drawing search, manufacturing data applications	Boundary and safeguarding analysis per Sections 7 and 8 before any deployment. Records: the documented analysis and decision, whatever the outcome

Three rules make the table operate. A use moves to a later stage only when its data analysis says so, and a use whose data reach is uncertain is analyzed at the higher stage until established otherwise. Embedded AI features are staged like any other tool, which means a feature arriving by update into a system that holds FCI is a Stage Two review on arrival, not a default. And every stage produces its records as the work happens, because the records are what convert the plan from good intentions into evidence.

Data class is the first axis of the plan, and it is not the only one. The second is operational consequence: what the tool's output does. A tool drafting website copy from public information and a tool recommending machining parameters or inspection acceptance from public equipment manuals sit in the same data stage and nowhere near the same risk. The plan therefore reads every proposed use on both axes, the sensitivity of the data the tool can reach, and the consequence of the output, from drafting assistance, through business decision support, to production, engineering, or quality decision support, to anything approaching acceptance authority or autonomous action. The approval level is set by the higher of the two: a use whose output influences production, quality, or acceptance receives Stage Two review depth and a defined human approval authority regardless of its data class, and acceptance or autonomous authority is a management decision with human review defined before deployment, not after. One page of discipline here is what separates an AI deployment plan from a data security plan presented as one.

4. Know the Shop's Data Before the Tools

The plan runs on one prerequisite: management has to know where the regulated information actually lives. In most small manufacturers, the honest answer is broader than the compliance documents say. CUI lives in the engineering repository and the drawing files, as expected, and also in the quoting system where technical requirements were pasted, in email threads with primes, in shared drives organized by job number, and sometimes in the shop floor systems that display work instructions. FCI is broader still, touching the ERP, the scheduling system, and most customer correspondence. The mapping exercise this section describes is not an enterprise data governance project; it is a focused pass, system by system, asking which regulated, proprietary, personal, privileged, and contract-restricted information categories each one holds or can expose, performed with the people who use the systems rather than from the network diagram.

The map does two jobs. Prospectively, it assigns every proposed AI use to its stage, because the stage of a use is the stage of the most sensitive data class in the systems it can reach. An AI feature inside the ERP is a Stage Two question in a shop whose ERP holds FCI, and a Stage Three question in a shop that pastes technical data into sales orders, and the map is how the company knows which shop it is. Retrospectively, the map joins the AI use inventory described in the companion paper to establish what the tools already in the building can currently reach, including the embedded features that arrived through updates and were never assigned to any stage at all. Where that retrospective look finds an AI capability already touching a class it was never reviewed for, the finding routes to the suspend outcome in Section 6, and the shop is better off for having found it before an assessor or an adjuster did.

5. Stage One: The Low-Risk Wins

Stage One exists because a plan that opens with prohibitions fails in a manufacturing business. The uses in this stage touch public information and approved low-sensitivity internal information only, no FCI, no CUI, no export-controlled technical data, and no personal, privileged, or contract-restricted information, and they deploy promptly: drafting for the website and marketing, capability statements built from already public information, administrative and internal writing that carries no technical content, scheduling and coordination help, and general research and drafting under organizational accounts. In most shops, this is the majority of what employees wanted AI for in the first place, and delivering it early does something no policy memo can, demonstrating that the company is prepared to approve AI use under governance rather than simply restrict it, which earns the credibility the later prohibitions depend on.

Stage One still runs under controls, and they are the inexpensive ones. Tools are used under organizational accounts on tiers whose terms have been read, not personal accounts, because the account tier is where data handling commitments live. The data rules from the AI use policy apply from the first day: the classes that never enter a Stage One tool or any other unapproved environment are named, and Stage One training teaches employees where Stage One ends, which is the single highest value training hour in the whole plan. That training states the plainest rule in this plan in words nobody can misread: CUI, FCI, and export-controlled technical data never enter a public generative AI chatbot or any tool running under a personal account, whatever the deadline, and the approved Stage One tools are the sanctioned way to get the help. This is a first-day rule rather than a Stage Three subtlety, because the most common AI violation in a defense shop is not an unvetted enterprise deployment; it is an employee pasting controlled information into a consumer chatbot to save an afternoon, and the rule has to reach that employee before the afternoon does. And the approved list starts here, in writing, so that the first records exist before the first hard question arrives. The management point of Stage One is momentum with evidence: the company banks visible value, employees experience the governed path as workable, and the discipline that Stages Two and Three depend on gets built where the stakes are low.

6. Stage Two: Business Systems and the Tool Review

Stage Two is where AI meets the systems the business runs on, quoting, planning, meetings, customer communication, and the ERP, and where the plan's tool review earns its place. The review is the vendor due diligence discipline compressed to what a small manufacturer needs to establish for each tool and each embedded feature before

approval. What the product does with content: whether prompts, files, and output are retained, whether content trains models, whether the training opt-out is a contractual commitment or a console setting someone must own, and whether human reviewers can see content. Who is actually behind the product: the model provider, the infrastructure, and the subprocessors, because the name on the invoice does not identify everyone who processes the data, and the NIST Generative AI Profile directs organizations to document and monitor exactly that chain. What the evidence proves: a SOC 2 report speaks to the controls and scope it describes, an ISO/IEC 27001 certificate to the vendor's security management system, an ISO/IEC 42001 certificate to its AI management system, and none of them establishes what the vendor may do with this customer's data, which lives in the contract and the data terms. And what the contract actually says: training prohibitions, retention and deletion, subprocessor disclosure and change notice, output ownership, and the liability position, with silence recorded as the finding it is.

Review area	**Acceptable answer**	**Conditional answer**	**Red flag answer**
Content retention	Defined, short, documented retention with contractual basis	Longer retention with a documented reason and deletion rights	Undefined retention, or retention rights reserved broadly
Model training use	Contractual commitment that customer content does not train models	Training excluded on the shop's tier, confirmed in writing	Content used for training by default, or opt-out limited to certain tiers
Human access	No human review, or review only on the customer's request	Review limited, logged, and disclosed	Undisclosed human review, or review by unidentified providers
Subprocessors	Named, located, and subject to change notification	Named with notification on material change	Unidentified, or changeable without notice

Stage Two divides at a line the table names and the review must respect: whether the tool merely operates near systems that hold FCI or actually processes, stores, or transmits it. An AI feature can run inside an ERP that contains FCI while being technically prevented from receiving it, and a feature that does receive FCI has crossed

into different territory: FAR 52.204-21 applies its fifteen basic safeguarding requirements to contractor systems that process, store, or transmit FCI, current CMMC Level 1 self-assessment is built on those same requirements, and the review for a Two-B deployment therefore adds the applicability analysis, confirmation that the service and its integrations sit inside the documented scope, and the recording of the service in the applicable system records. The distinction is established by the data map and the tool's connectors, not by the vendor's marketing, and a tool whose FCI reach cannot be established is reviewed as Two-B until it can.

Stage Two also owns the embedded feature problem for the whole plan. AI features now arrive inside approved software through routine updates, often enabled by default, and no purchase order announces them. The control is a standing posture rather than a gate: new AI features default to disabled or restricted until reviewed, someone owns the vendor release notes and administrative consoles for the systems that matter, and every feature decision routes into the inventory. Each Stage Two review ends in one of five outcomes: approve, approve with restrictions, pilot with a review date, reject with reasons recorded, or suspend when an approved product materially changes, and each ends in a review file and a one-page approval decision record naming the use, the data class, the restrictions, the owner, and the next review date. The file matters as much as the decision it records, because it is the shop's answer, in writing, when a prime's questionnaire or the insurance application asks how AI vendors are overseen.

7. Stage Three: The CUI Question

Stage Three is the section most readers of this paper will turn to first, because the highest value uses in a defense shop, analyzing engineering documents, searching technical drawings, working with manufacturing data, are exactly the uses that touch CUI. The analysis has to be stated without softening. An AI service that would process, store, transmit, summarize, index, or generate output from CUI has entered the safeguarding analysis under NIST SP 800-171 and DFARS 252.204-7012. Structurally, that AI vendor is an external service provider whose service happens to be inference, and the analysis then divides by how the service is delivered. A cloud AI service processing CUI is evaluated as a cloud service provider: under DFARS 252.204-7012, the specific product or service offering, not merely an underlying platform, must meet security requirements equivalent to those established for the FedRAMP Moderate baseline and comply with the clause's incident reporting and forensic support provisions. A terminology note applies as of mid 2026: the FedRAMP program's Consolidated Rules for 2026, finalized in June, retired the Low, Moderate, and High labels in favor of certification Classes A through D, with the former Moderate baseline mapping to Class C and authorized services now designated FedRAMP Certified. The DFARS clause's Moderate baseline language governs until the clause is amended, so the

review applies the clause's standard while recording a provider's status in the program's current terminology. A provider that is not a cloud service but processes CUI is brought inside the company's own compliance perimeter: documented in the SSP and included in the assessment scope. In either case, the questions are the ones that apply to any ESP: what covered information it can access, whether it meets the applicable requirements, and what its presence does to the assessment boundary and the SSP. An AI service inside the boundary must be described, safeguarded, and assessable there; a service outside the boundary must be shown to receive no covered information, and the drawing of that line is precisely the scoping discipline defense contractors already know from the rest of their compliance lives.

The honest current state is that most commercial cloud AI services cannot make the commitments this analysis requires. Consumer and standard business tiers rarely offer the contractual data handling, the flowdown acceptance, or the incident reporting posture that covered information demands, and a shop cannot cure that gap with a configuration setting. The realistic Stage Three options are therefore three. The first is a specialized environment: the government-focused cloud offerings built to hold covered information, where AI capabilities operate inside an environment designed for the data, evaluated with the same ESP discipline as any other. The second is local operation: models run on the shop's own systems, keeping covered data within the contractor-controlled environment. Local operation solves the external transmission problem and does not eliminate the safeguarding analysis, because the model, its runtime and dependencies, its administrative accounts, its logs, and its data flows become components of the environment, added to the SSP, the asset inventory, the security architecture, and the assessment evidence, with the capability and upkeep costs weighed honestly beside that documentation work. The third is the decision not to deploy, which this plan treats as a fully legitimate outcome rather than a failure, because a documented decision that the current options do not meet the shop's obligations is itself governance operating.

Management should price one more line before choosing among the three: the System Security Plan work that any Stage Three deployment carries. AI integrated into the CUI environment is a material change to the system the SSP describes, and the documentation follows the change. The boundary and data flow diagrams gain the inference path, from prompt through retrieval to output and logs. The component inventory gains the model, the runtime, and their dependencies. The implementation statements for the affected requirements are revisited, among them access control over the model and its interfaces, audit content that now includes prompts and outputs, configuration and change management extended to model and version changes, and incident response extended to AI behavior. Derived data such as embeddings and indexes built from CUI is handled as CUI wherever it lands, and for a cloud service the

shared responsibility documentation joins the file. The change also bears on what the company has already asserted: a materially changed environment is reviewed against the current assessment and affirmation basis before the next score is entered or relied on. None of this argues against Stage Three; it prices Stage Three honestly, and a shop that has priced it can decide, while a shop that has not will discover the price during an assessment.

The audit and accountability requirements deserve their own line in that pricing, because they are where AI deployments most often fail quietly. Once an AI capability is a component of the CUI environment, the organization defines auditable events for it and retains the records: the prompts and files submitted, the repositories and documents a retrieval feature actually reached, the outputs produced, the identity of the user behind each interaction, the model and version in service, administrative and configuration changes, and any action the capability takes on other systems, because an AI feature that acts is an actor whose actions must be traceable like any user's. The requirement does not extend to the model's internal computation, which no framework demands and no service can meaningfully produce; it extends to what entered, what was reached, what came out, who was responsible, and what changed, in records the company can retain and review. A capability that cannot produce that account of itself is not ready for the environment, whatever its other credentials, and for locally operated models the shop owns that logging configuration itself, which belongs in the same SSP work priced above.

Whichever outcome the analysis reaches, it is written down: the systems and data considered, the options evaluated, the requirements applied, and the decision with its owner and date. The record matters beyond good order. The company's representations, in the SSP, in SPRS, to primes, describe a controlled environment, and the documented Stage Three analysis is the company's evidence that AI entered that environment only by decision, or did not enter it at all. In an assessment or an incident, that record is what distinguishes a company that governed its environment from one that assumed the question would never be asked.

8. Export-Controlled Technical Data

Export-controlled technical data raises a separate analysis from CUI, and this section's job is to put it on the management agenda rather than to resolve it. Technical data controlled under the International Traffic in Arms Regulations or the Export Administration Regulations is restricted in who may access it and where it may go, and transmitting it to a cloud AI service raises export questions that ordinary security certifications do not answer: where processing occurs, who can access decrypted content, who controls the encryption and decryption keys, and whether the arrangement

satisfies an applicable authorization, exemption, or exception. The provider's architecture is directly relevant here, because the regulatory provisions treating properly secured end-to-end encrypted transfers as something other than exports turn on whether the provider can access the content, and an AI service generally must decrypt a prompt or file to process it. The supply chain mapping from Stage Two does double duty here, because the model providers and subprocessors behind an application are part of the access analysis, and a vendor that cannot identify where content is processed has answered the export question in the negative.

The plain guidance for a small manufacturer is short. The data map from Section 4 identifies where export-controlled technical data lives. No AI tool touches those systems or that data until export counsel has reviewed the specific arrangement, and the review is documented like every other decision in this plan. The shops most at risk here are not the ones that ask the question and get a hard answer; they are the ones where an engineer, trying to be efficient, summarizes a controlled drawing package with a consumer AI tool on a Tuesday afternoon, which is a Stage One training topic and an access control matter long before it is a legal one.

9. The Floor Rules at Every Stage

Beneath the stages, a small set of program elements applies from the first day and does not change as deployments advance. The AI use policy makes the eight decisions described in the companion paper, with the data rules leading in a defense shop: the classes that may never enter an unapproved AI tool, with FCI, CUI, and export-controlled technical data at the top, together with the specific environments, account tiers, contractual conditions, and approval requirements under which each class may be processed at all, which is the policy expression of the stages themselves. Training and acknowledgment run with the Stage One rollout and repeat for new hires, because the policy nobody was trained on is an unenforced policy, and in this environment an unenforced policy sits next to attested representations. The embedded feature watch runs continuously, with a named owner. And the inventory, the approved list, the review files, and the decision records are kept current as the work happens, not reconstructed when a question arrives.

The floor rules also govern what the output does, because in a manufacturing business the consequence axis from Section 3 has to be operated, not merely acknowledged. Each approved tool carries a defined intended use and named prohibited uses. Output that influences production, engineering, quality, or customers passes a defined human review before it is relied on, with the approval authority named. Fitness is tested proportionately to consequence before a tool graduates from pilot. A vendor's change of model or material feature is a standing reassessment trigger, logged when it occurs.

AI-related errors route into the quality system the shop already runs, through the same corrective action discipline as any other nonconformance. And generated output is classified according to its content, with a conservative operating rule stated in writing: output produced from CUI or export-controlled technical data is handled at the source's protection level until an authorized reviewer determines otherwise.

The return on the floor rules is representational, and for this audience that is the argument that matters. A shop running this plan can answer, accurately and from its files, the questions it is already being asked: the prime's questionnaire about AI use and oversight, the cyber insurance application's questions about AI governance, policy, and training records, and any future assessment conversation about what the environment contains. Each stage of the plan produces the records from which those answers are drawn. A shop that deploys AI without the records is not more advanced than one that waited; it is merely making representations it cannot support, which is the oldest failure mode in this industry and the one every reader of this paper already knows how to name.

10. The Plan at a 50-Person Shop

This closing section shows the plan in practice, at a 50-person precision manufacturer holding defense contracts with flowdown clauses and a cyber insurance renewal every fall. In the first quarter, management holds the Step Zero meeting and names the owner. The data map takes three weeks of part-time effort and finds what these maps usually find: CUI in the engineering repository and the job folders as expected, and also in the quoting system and in years of email with two primes. Stage One deploys in the same quarter: the office productivity assistant on the enterprise tier under organizational accounts with training use excluded, deployed initially as a standalone drafting interface with its email, drive, meeting, and ERP connectors disabled, because the data map just found CUI in the quoting system and in email, and each connector will be enabled only as its repositories are cleared or excluded by technical control. Marketing and administrative drafting are encouraged from the first week, and the policy rolls out with an hour of training whose central topic is where Stage One ends. The approved list exists in writing by the end of the quarter, and so do the first acknowledgment records.

In the second quarter, Stage Two runs four reviews. The ERP vendor's new forecasting feature is piloted for ninety days against internal planning data, with its AI addendum read for the first time. The meeting assistant is approved with restrictions, limited to meetings designated as free of FCI, CUI, export-controlled technical data, and personnel or privileged content, because internal is a location rather than a data classification, and because recording customer conversations engages consent and confidentiality questions the shop has not cleared. The customer service drafting tool is approved with

a defined human review boundary, because its output reaches customers. And the review of the environment against the data map finds that a document search feature, enabled by a vendor update in the spring, has been indexing a shared drive that holds FCI; the feature is suspended, the finding is recorded, and the shop is glad to have been the one to find it. In the third quarter, the Stage Three analysis is documented with the shop's compliance advisor: the engineering document analysis the production manager wants is evaluated against the safeguarding requirements, the commercial options cannot meet them, the specialized environments are priced with the SSP and assessment documentation counted in the price, and the decision is deferred to next year's planning, and the documented decision not to deploy closes the analysis. When the insurance renewal arrives in the fall, the AI questions on the application are answered in an afternoon, from the files.

At no point in those three quarters did the shop stop being a manufacturer to become a compliance project. The plan ran beside production, in meetings the company was mostly holding anyway, and that is the only way it ever actually runs. The result is a manufacturing business rather than an AI transformation project: one that adopted the AI that served the operation, declined what its obligations required, and holds the records to demonstrate both decisions.

About the Author

David W. Koran is a CyberAB Registered Practitioner Advanced (RPA) and holds ISO/IEC 27001 certifications as an auditor and in implementation. He is the founder of a consulting practice serving Defense Industrial Base contractors and their legal counsel, focused on readiness, enablement, and implementation, and advises organizations on AI governance. He is an Associate Member of the American Bar Association Section of Public Contract Law and the author of The CMMC Decision. He can be reached at dkoran@davidkoran.com and (802) 335-2662.

References

A version note on NIST SP 800-171: Revision 3 is the current NIST publication, while DFARS 252.204-7012 applies the version in effect for the applicable solicitation or as authorized by the contracting officer, and current CMMC materials remain based on Revision 2. Contractors determine the applicable revision per contractual obligation, which is why both are cited below.

National Institute of Standards and Technology, Special Publication 800-171 Revision 2, Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations.

National Institute of Standards and Technology, Special Publication 800-171 Revision 3, May 2024.

Defense Federal Acquisition Regulation Supplement 252.204-7012, Safeguarding Covered Defense Information and Cyber Incident Reporting.

Federal Risk and Authorization Management Program, Consolidated Rules for 2026 and certification class transition, June 2026.

National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1, July 2024.

National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, January 2023.

International Organization for Standardization, ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system.

International Traffic in Arms Regulations, 22 CFR Subchapter M.

Export Administration Regulations, 15 CFR Subchapter C.

David W. Koran, How to Build an AI Governance Program: A Records-Based Guide for Small and Midsize Organizations That Use AI, August 2026.