How to Build an AI Governance Program

A Records-Based Guide for Small and Midsize Organizations That Use AI

David W. Koran

CyberAB Registered Practitioner Advanced

ISO/IEC 27001 Auditor and Implementer

August 2026

Summary

An AI governance program is the set of structures, decisions, and records through which an organization directs and controls its use of artificial intelligence. This paper describes how to build one, in order, for a small or midsize organization that deploys and uses AI rather than developing it, which is the situation of most companies now facing AI questions from customers, insurance carriers, contract clauses, and state law.

The program has five components: an AI use inventory, an AI use policy, alignment with a governance framework, oversight of AI vendors and embedded AI features, and sustainment. Each component is defined in this paper by the records it produces, because records are what every statute, framework, customer, and carrier ultimately asks for. A governance program that cannot produce its records on request has not finished being built, however good its intentions.

The paper opens with the question executives ask first: why build now, when the rules keep moving? The answer developed in Section 1 is that the instability is concentrated in statutes while the frameworks underneath them have held still, that one state has written framework compliance into its statutory defense structure, and that organizations are making binding representations about their AI use today whether or not a program exists behind them. The remaining sections build the five components in sequence, and the paper closes with a worked profile of the program at a fifty-person manufacturer and a one-page build sequence that can be handed to whoever will own the work.

1. Why Build Now, Against Rules That Keep Moving

Between 2024 and 2026, the legal ground under AI governance moved repeatedly. Colorado enacted the first comprehensive state AI statute in 2024, amended it, paused its enforcement, and in May 2026, before it ever operated, repealed and reenacted it as an automated decision-making transparency law, SB 26-189, whose obligations begin January 1, 2027. The European Union, having adopted the Artificial Intelligence Act in 2024, revised its principal high-risk application dates in the June 2026 Digital Omnibus amendments, moving the Annex III high-risk requirements to December 2, 2027 and the product-related Annex I requirements to August 2, 2028. A December 2025 federal executive order directed agencies to pursue preemption of state AI laws, and that federal preemption initiative remains legally and politically unresolved. That record supports a conclusion that sounds reasonable in any budget meeting: building a compliance

program against standards still being rewritten looks like wasted money, and the prudent move looks like waiting for the rules to settle.

The conclusion fails for two reasons. The first is that waiting is not neutral. While an organization waits, its employees adopt AI tools, its vendors embed AI features in products already in service, and its AI footprint grows without an inventory, a policy, or a record. More consequentially, the organization keeps making representations during the waiting period. It answers AI questions on cyber insurance renewals, signs contracts containing AI clauses, and responds to customer due diligence questionnaires. Those answers are commitments, and commitments made without supporting records are the oldest failure mode in compliance. Defense contractors will recognize the structure: a score entered in a government system without the evidence behind it is not a neutral act, and neither is an insurance application answer about AI governance that nothing in the organization can substantiate.

The second reason is that the instability is not evenly distributed. It is concentrated in statutes, while the frameworks underneath them have held still. The NIST AI Risk Management Framework, published in January 2023, has not changed while Colorado legislated, repealed, and replaced; neither has ISO/IEC 42001, the certifiable AI management system standard published in December 2023. Texas made the relationship concrete, and the mechanism deserves precise statement. The Texas Responsible Artificial Intelligence Governance Act, in force since January 1, 2026 and a major cross-sector state AI statute currently operating, builds framework compliance into its defense structure: a person may avoid liability where, among other conditions, a violation is discovered and addressed through mechanisms such as good faith testing, audits, or internal review, and the person substantially complies with the NIST AI Risk Management Framework, including its Generative AI Profile, or another nationally or internationally recognized AI risk management framework. That is narrower than a blanket safe harbor, and it is more instructive for it, because what the statute protects is a documented, operating governance process rather than a certificate of alignment. Colorado's original statute contained its own framework-based defense; when the law was repealed and reenacted, the defense did not carry forward. The lesson runs one way: an organization that governed to the Colorado statute spent money on a regime that no longer exists, while an organization that governed to the framework holds records that the Texas statute rewards, that procurement questionnaires recognize, and that every plausible future regime will draw on.

Adoption research sharpens the timing. Surveys through 2026 consistently find the same shape regardless of publisher: AI deployment has far outrun formal governance, with most organizations reporting AI in use while only a minority report a formal governance authority over it, and analyst projections converge on formal AI governance programs becoming the enterprise norm by the end of 2026. The organizations closing

that gap are not doing so because a regulator arrived. They are doing so because the enforcement mechanisms that already reach them- contract clauses, prime contractor flowdowns, underwriting questions, and existing data protection obligations such as the safeguarding of Controlled Unclassified Information under NIST SP 800-171 and DFARS 252.204-7012, have started asking AI questions in writing. The practical conclusion of this section, and the premise of the rest of the paper, is that the durable investment is the program and its records, built to stable frameworks, with statute-specific analysis layered on top as thin and replaceable material.

2. The Architecture in One View

The program has five components, and they arrive in a deliberate order. The inventory establishes what AI the organization actually uses. The policy converts that knowledge into decisions about what is permitted, prohibited, and subject to approval. Framework alignment structures the risk work and produces the vocabulary counterparties expect. Vendor oversight examines the AI the organization buys, which for most small and midsize companies is all of it. Sustainment keeps the other four accurate as tools, vendors, and obligations move. Later sections build each component; the table below states what each one is and the records it produces, because the records are the working definition of done.

Component	**Records it produces**
1. AI use inventory	The inventory register: each AI system, tool, and embedded feature; its purpose and intended use; the data it can reach; the accounts and terms it runs under; the vendor behind it; whose residents it can affect; and the register's review date.
2. AI use policy	The policy itself; training and acknowledgment logs; approval decisions for approval tier uses; the exception register; policy review minutes.
3. Framework alignment	The gap analysis against the NIST AI RMF and, where pursued, ISO/IEC 42001; documented risk decisions; intended use documentation; impact assessment records; the Statement of Applicability where the ISO standard is pursued.
4. Vendor and embedded AI oversight	A review file per vendor and per embedded feature: the terms and questionnaire responses relied on, the configuration decisions, the governing contract clauses, the decision with its owner and date, and the next review date.

Component	Records it produces
5. Sustainment	Review minutes on the defined cadence; the change log connecting tool, vendor, and regulatory changes to updates in the other four record sets.

One scaling note governs everything that follows. The program described here is the deployer and user version, built for an organization that buys and uses AI rather than developing it. A company that trains models or integrates AI into products it sells carries additional obligations, principally in testing, documentation, and the representations it makes to its own customers, and Section 8 notes where the divergence begins. For the majority of small and midsize organizations, the deployer-and-user program is the honest scope, and it is a bounded, buildable project rather than an enterprise transformation.

One step precedes component one, and skipping it is the most common cause of stalled programs. Before the inventory begins, management establishes the mandate: an executive sponsor, a named program owner with genuine decision authority, the scope of the organization being governed, the escalation path for decisions above the owner's authority, and the reporting cadence to leadership. This is not a sixth component; it is the authorization the five components run on, and it is the reason the build sequence in Section 8 begins with management ownership rather than with technology.

3. Component One: The AI Use Inventory

The inventory comes first because every other component depends on it. A policy written before the inventory bans tools people depend on and permits categories nobody uses, forfeiting its credibility on publication day. A framework gap analysis without an inventory assesses a fiction. Vendor oversight cannot begin until the organization knows which vendors are actually supplying AI. The inventory is also where most organizations receive their first genuine surprise, because the inventory that emerges from evidence is reliably larger than the one management would have written from memory.

The surprise has two sources, and the inventory must cover both. The first is employee adoption: people under deadline pressure discover that a consumer AI tool drafts, summarizes, or codes faster than they can, sign up with personal accounts, and begin using it for work. Nothing about this is malicious; it is the same dynamic that produced shadow IT a decade earlier, and it means organizational data is flowing to services the organization has never evaluated, under consumer terms nobody has read, through accounts it cannot see or suspend. The second source is quieter and now arguably

larger: vendors adding AI features to software the organization already runs. The office suite gains an assistant, the customer relationship system gains generative drafting, the meeting platform gains transcription, often enabled by default in a routine update. No procurement decision occurred because nothing was procured. An inventory built from the purchasing records will miss this entire category, which is why the inventory is built from the environment instead.

Building from the environment uses evidence the organization already generates. Web filtering, DNS, and firewall logs show traffic to AI service domains; the question is rarely whether the evidence exists but whether anyone has been asked to look for that category of destination. Identity records, single sign-on logs, and OAuth grants reveal AI services connected to organizational accounts, and browser extension inventories reveal the AI riding inside the browser. Vendor administrative consoles and release notes reveal the embedded features and their default states. The final source outproduces all of the others: structured interviews. People will describe what they use when the question is framed as inventory rather than enforcement, and the practitioner experience across every compliance framework is the same here as elsewhere: the people who run the operation know how it actually runs.

Each inventory entry records the system or feature, its business purpose and intended use, the data it can reach, the accounts and terms it operates under, the vendor supplying it, and one field whose importance has grown with the state statutes: whose residents the system can affect. State AI laws generally attach to the people a system touches rather than to the deployer's location, so a hiring tool that screens an Illinois applicant or a chatbot serving Texas consumers sits inside those regimes wherever the company is headquartered. Jurisdictional exposure is a property of the user population, and the inventory is the only place it can be recorded. Records produced: the inventory register, with an owner and a review date, which becomes the foundation document for everything that follows.

4. Component Two: The Decisions and the AI Use Policy

An AI use policy is a set of decisions, not a document about intentions. The distinction matters because the most common AI policy in circulation is a downloaded template with its decisions unmade: which tools are approved is left generic, which data classes are excluded is left abstract, and who approves new uses is left to an unnamed committee. A template in that condition instructs no one, and the gap between what it says and what people do becomes a finding the first time anyone compares them. The policy earns its place in the program by making eight specific decisions.

First, scope and definitions: what counts as an AI tool for purposes of the policy, including consumer chatbots, embedded features, browser extensions, and AI-enabled services, and who is covered, meaning employees, contractors, and anyone acting with organizational data or accounts. Second, the three-tier use decision: uses permitted outright, uses prohibited outright, and uses requiring approval. Three tiers work where binary rules fail, because most AI use is neither clearly safe nor clearly unacceptable, and the approval tier doubles as the organization's ongoing intelligence about what its people actually need. Third, the data rules: the classes of information that may never enter an AI tool regardless of tier, keyed to the organization's actual obligations. For a defense contractor, that list begins with Controlled Unclassified Information; for most organizations it includes personal information and customer data held under confidentiality terms. This clause protects the representations the organization has already made to outside parties. Fourth, the approved tools and accounts: the sanctioned path, under organizational accounts and reviewed terms rather than personal ones. The approved list is what makes the prohibitions enforceable in good conscience, because the announcement that some uses are ending has to arrive together with the sanctioned way to keep the benefit.

Fifth, embedded and vendor AI handling: who reviews vendor AI feature announcements and administrative settings, which features default to off pending review, and how each review outcome reaches the inventory. Sixth, human oversight and output responsibility: the principle that AI-assisted work product remains the responsibility of the person who produces it, the review that applies before AI-assisted output reaches customers, code, or decisions, and any disclosure the organization requires about AI involvement. Seventh, training and acknowledgment: how the policy is taught, when acknowledgment is collected, and how new hires are covered, because a policy nobody was trained on is an unenforced policy. Eighth, enforcement, exceptions, and review: what happens on violation, how exceptions are requested and recorded rather than improvised, and the review cycle, which in this domain means at least semiannually because the tools and the obligations both move faster than an annual cycle.

Ownership of these decisions belongs to management, not to the IT function. The policy allocates risk the organization is accepting, restricts how regulated and contractual information is handled, and supports representations made to customers, carriers, and regulators; those decisions belong to leadership, with the IT and security functions implementing and monitoring what leadership has decided. Records produced: the policy, the acknowledgment logs, the approval decisions, and the exception register, which together are usually the first artifacts an outside party requests and the clearest early evidence that the program operates.

5. Component Three: Framework Alignment

The framework gives the program its structure and its vocabulary, and for a United States organization, the framework is the NIST AI Risk Management Framework. Published as AI RMF 1.0 in January 2023, it organizes AI risk management into four functions, Govern, Map, Measure, and Manage, elaborated through nineteen categories and seventy-two subcategories that describe outcomes rather than prescribing methods. The framework is voluntary, and its weight is referential: Texas builds substantial compliance with it into the statutory defense structure described in Section 1, procurement and due diligence questionnaires name it as the expected answer, and published crosswalks connect it to ISO/IEC 42001 and the EU AI Act so that work performed against it transfers rather than duplicates.

For a deployer and user organization, the four functions map onto work already described in this paper, which is the point: framework alignment is not a sixth activity bolted onto the other five but the structure that organizes them. Govern holds the inventory, the policy, assigned accountability, and third-party oversight; it is the function where gaps concentrate in organizations that adopted tools before adopting governance. Map documents each significant system's context: what it is for, who it affects, the data involved, and what could foreseeably go wrong, producing the intended use documentation that an intent-based statute such as the Texas Act rewards. Measure, at deployer scale, means evaluating what a deployer can genuinely evaluate: vendor claims and documentation, contract terms, outputs examined against intended use, and the indicators that reveal drift; it does not mean model testing laboratories, and pretending otherwise is how small organizations talk themselves out of the framework entirely. Manage turns measurement into operation: prioritizing risks against tolerance, applying treatments, monitoring systems in service, and deciding when a system should be constrained or retired. Because generative AI is precisely the AI most organizations use, the Generative AI Profile, NIST AI 600-1, published in July 2024, applies almost everywhere: it identifies twelve risk categories unique to or amplified by generative systems and maps deployer-level actions to the four functions.

ISO/IEC 42001 is the institutionalization path rather than a competing framework. Published in December 2023, it is the first certifiable standard for an artificial intelligence management system, with its requirements in Clauses 4 through 10 on the same harmonized structure as ISO/IEC 27001 and a reference set of thirty-eight controls in nine Annex A objectives, selected through a Statement of Applicability. Its two genuine additions beyond an existing information security management system are the AI-specific control set and the AI system impact assessment established at Clause 6.1.4 and performed at Clause 8.4, which asks what the organization's AI could do to people and has no information security equivalent. An organization already operating

an ISO 27001 management system extends what it has, using the same document control, internal audit, and management review, rather than building a parallel bureaucracy.

The certification decision deserves an honest rule rather than a sales conversation. Certification is worth pursuing when customers or markets will ask for the certificate itself, which today is most common for technology vendors, AI-enabled service providers, and organizations selling into large enterprises or European buyers, where the standard has begun appearing in procurement requirements. For most other organizations, alignment without certification is the sensible position: build the management discipline the standard describes, hold the records, and certify later if the market demands it. The management system is the durable asset; the certificate is a claim about it that can be added when it earns its cost. Records produced: the gap analysis, documented risk decisions, intended use documentation, impact assessment records where the ISO path is taken, and the Statement of Applicability, all of which double as the answers to the framework questions now arriving in questionnaires.

6. Component Four: AI Vendor and Embedded AI Oversight

Most organizations do not build AI; they buy it, which means most AI risk is vendor risk, running through terms the organization did not draft and behavior it does not control. Ordinary vendor review asks whether the vendor will protect the data. AI vendor review additionally asks what the vendor will do with it, and the review examines six areas for each vendor and each embedded feature.

The first is the data handling and training terms: whether customer content is retained and for how long, whether it is used to train or improve models, whether an opt-out exists and is actually configured, and whether human reviewers see content. This clause set determines whether using the tool is compatible with the organization's confidentiality obligations at all. The second is retention, deletion, and subprocessors: where the data physically goes, including the underlying model providers, because AI services are frequently assembled from other AI services and the organization's data travels the whole chain. The third is security posture, examined with precision about what each artifact proves: a SOC 2 report and an ISO 27001 certificate speak to information security, an ISO/IEC 42001 certificate speaks to the vendor's AI management system, and none of them answers the data handling questions, which live in the contract. The fourth is the vendor's representations about the AI itself, its intended use, limitations, and testing, because those representations are what the organization will cite when its own use of the tool is questioned. The fifth is contract

allocation: training use restrictions, confidentiality scope, ownership of AI output, indemnification for the vendor's AI behavior, notice obligations when the vendor changes models or features, and audit rights, with the observation that standard contract templates predate these questions, so silence in the contract is itself a finding. The sixth is the change path: assigned ownership for reviewing AI feature announcements and administrative settings, a default posture for new features pending review, and the route by which each change reaches the inventory. Without the sixth area, the review describes the vendor's product only as it existed on the day the review was performed, and the vendor's next release quietly retires it.

Defense Industrial Base organizations carry a sharper version of every question above. An AI service that touches Controlled Unclassified Information enters the safeguarding analysis under NIST SP 800-171 and DFARS 252.204-7012, and potentially the assessment scope; structurally, an AI vendor is an external service provider whose service happens to be inference, and the analysis that applies to managed service and external service providers applies to it. Records produced: a review file per vendor and per embedded feature containing the terms and responses relied on, the configuration decisions, the governing clauses, the decision with its owner and date, and the next review date. The file earns its cost twice, once in the decision it informs and again every time a counterparty asks the question and the answer is already on hand.

7. Component Five: Sustainment

AI governance decays faster than most compliance programs, for reasons built into its subject. The tools change monthly, vendors ship AI features on their own schedules, employees adopt new services continuously, and the statutes themselves are being enacted, amended, and repealed in real time. A program built in one quarter and left untouched is materially out of date within two more, which is why sustainment is a component of the program rather than an afterthought.

Sustainment is a calendar with owners. The policy is reviewed at least semiannually, with the approved tool list checked for accuracy, because an outdated approved list undermines the whole document. The regulatory picture is reviewed quarterly, tracking the scheduled changes, the Colorado successor statute and the other 2027 arrivals, the European dates, and the federal preemption question, and adjusting the thin statutory layer while the framework base stays put. Event-driven reviews fire when a major tool, vendor feature, or obligation changes, routed through the change path that vendor oversight established. Each review updates the registers and leaves minutes, and the whole rhythm integrates with the management cycles the organization already runs: the information security management system, the System Security Plan maintenance, the insurance renewal, rather than standing beside them as a parallel bureaucracy. Records

produced: review minutes on the cadence, the change log, and registers whose review dates are current, which is the single fastest thing an outside reviewer checks.

8. The Program at Working Scale

What does all of this amount to at a fifty-person manufacturer that deploys and uses AI, holds government or commercial contracts with flowdown clauses, and answers a cyber insurance application every year? The inventory typically lands between fifteen and forty entries once embedded features are counted, and building it is two to four weeks of part-time effort: the log review, the identity and extension sweep, the vendor console check, and a dozen interviews. The policy is a working session with management to make the eight decisions, a draft written in the organization's own language, and a rollout paired with the approved path; from inventory completion, it is buildable in weeks. Framework alignment is a gap walk through the AI RMF subcategories at deployer depth with the Generative AI Profile applied, producing a findings list and the intended use documentation; the ISO 42001 question is answered with the certification rule from Section 5, which for this profile usually means alignment now and certification only if a customer asks. Vendor oversight opens with review files for the five to ten vendors and features that touch meaningful data, and the change path assigns one named person to watch release notes and settings. Sustainment is the calendar: semiannual policy review, quarterly regulatory review, and the change log. None of this requires a compliance department; it requires an owner with authority, a defensible sequence, and the discipline to produce records as the work happens rather than reconstructing them later.

An organization that develops or integrates AI diverges upward from this baseline, principally in Measure, where testing and documentation obligations become real engineering work, in the representations it makes to its own customers, which become the subject of someone else's vendor due diligence, and in the impact assessment, which deepens from a deployer's context review into a development lifecycle discipline. The five component architecture holds at that scale as well, with the depth of each component increasing rather than the structure changing.

The build sequence, in one list an owner can carry: confirm management ownership and the program's scope; build the inventory from the environment, including embedded features and the resident reach field; make the eight policy decisions and roll the policy out together with the approved path, collecting acknowledgments; walk the AI RMF gap analysis at deployer depth with the Generative AI Profile, documenting intended use and risk decisions; decide the ISO/IEC 42001 question with the certification rule; open the vendor review files and assign the change path; set the sustainment calendar and hold the first review on schedule. At the end of that sequence, the organization holds a

working program and, more to the point, holds its records, which is the position from which every AI question now arriving, from a statute, a customer, a carrier, or its own board, can be answered accurately and without improvisation.

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

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

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

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

Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act), as amended.

Texas Responsible Artificial Intelligence Governance Act, House Bill 149, 89th Legislature, effective January 1, 2026.

Colorado Senate Bill 26-189, Artificial Intelligence Consumer Transparency, 2026.

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

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