



Regulated AI: From Tools to Systems

A practitioner's manual for operationalizing AI compliance infrastructure in regulated organizations.

A Dang's Solutions Publication

By Dan Gonzalez, Founder, Dang's Solutions

CONTENTS

About Dang's Solutions

Who This Guide Is For

How to Use This Guide

Part 1 · The Case for AI Infrastructure in Regulated Organizations

- 1.1 The Ungoverned Environment
 - 1.2 The Cost of Inaction vs. the Cost of Action
 - 1.3 The ROI Framework
-

Part 2 · What "AI-Ready" Actually Means for a Regulated Organization

- 2.1 Data Readiness
 - 2.2 Document Readiness
 - 2.3 Governance Readiness
 - 2.4 Infrastructure Readiness
 - 2.5 From Assessment to Plan: The Four Phases
 - 2.6 Building Your Readiness Plan
-

Part 3 · The Orchestration Layer

- 3.1 Why Data Sources Are Not Ready Today
- 3.2 The Compliance Integrator Role
- 3.3 The Standard Your Peers Are Moving Toward
- 3.4 What Comes Next

About Dang's Solutions

Dang's Solutions provides compliance consulting and AI governance for regulated organizations. The work started because I was often the human doing the manual review, and I wanted the mechanical parts of that review to run without me so my time could go to the findings. Dang's Solutions encodes practitioner judgment as compliance skills: self-contained, auditable procedures that a model can execute and a practitioner can verify.

Who This Guide Is For

This guide is for organizations and individuals evaluating or deploying generative AI. It is a practitioner's manual for operationalizing AI compliance infrastructure, written for the person who owns the regulatory obligation rather than the person who administers the network.

It assumes you already know your industry. You know what a risk assessment is, what a vendor contract is supposed to say, and what happens when a regulator asks a question you cannot answer. What this guide adds is the AI governance layer on top of that expertise.

How to Use This Guide

Part 1 establishes the utility of deploying AI tools, in numbers rather than adjectives. Part 2 is the working section: four readiness dimensions, each with a scored instrument you can run on your own organization, a scorecard that consolidates the four results, and the four-phase remediation path those results sequence into. Part 3 addresses the interaction between the AI and the human overseeing and directing the activity.

Every section contains at least one framework, scored instrument, or decision tool you can use.

The Case for AI Infrastructure in Regulated Organizations

AI is probably already inside your organization. It may have arrived through individual accounts, through features your vendors switched on, or through work product nobody labeled. Governance may have arrived later, or not at all. This part accounts for what that gap costs you, what evaluating and correcting any deficiencies looks like, and what completing the work returns.

1.1 The Ungoverned Environment

Most estimates put individual AI usage among professionals somewhere between a third and three-quarters, depending on the industry and how you define "using AI." The exact number matters less than the gap. At most organizations, formal training and governance have not kept pace, and a significant share of organizations provide no formal AI training at all.

The gap shows up in many forms. Three of them recur across industries often enough to plan around, and they share a root: the organization has lost track of its own environment. Only the first of the three usually gets a name, shadow AI, and the name is narrower than the problem it describes.

Failure mode 1: Regulated data entered into public AI tools.

Data flows through individual accounts under consumer terms rather than a Business Associate Agreement or Data Processing Agreement. The immediate cost is a complete loss of visibility. Because the disclosure is unlogged, the organization cannot account for what data was processed or when. You carry the regulatory risk without knowing its scope.

IN PRACTICE · FAILURE MODE 1

A staff member pastes a document into a consumer AI tool to summarize it. The document contains Protected Health Information, a client confidence, or material non-public information. Months later, you are asked to prove exactly who has had access to that data. You cannot, because your systems have no record that this disclosure ever happened.

Cited incident. In early 2023, three engineers at Samsung Semiconductor pasted proprietary source code, equipment-defect data, and internal meeting transcripts into ChatGPT's consumer interface within twenty days of the tool being permitted for work use. Because the consumer product retained inputs for model training, Samsung's confidential data was irrecoverably transferred to a third party's servers with no audit trail, no contractual coverage, and no mechanism for retrieval or deletion. Samsung subsequently banned all external generative AI tools company-wide. (Samsung internal disclosure, widely reported March–April 2023; Samsung confirmed the incidents and the ban in public statements.)

Failure mode 2: AI-generated work product with no audit trail.

Work product carries AI-generated content that nobody recorded as such. The output may be entirely correct. The exposure is that you cannot reconstruct how a conclusion was reached once somebody challenges it.

IN PRACTICE · FAILURE MODE 2

An analysis goes out with an AI-drafted section in it. Nobody logged which tool produced it, what it was given as input, or who verified the output. A year later a regulator asks how you reached that conclusion. Your defensibility depended on a record that was never created.

Cited incident. In *Mata v. Avianca, Inc.*, No. 22-cv-1461 (S.D.N.Y. June 22, 2023), attorneys submitted a brief citing six judicial opinions that did not exist. The opinions had been generated by ChatGPT; no one verified them against any legal database before filing. When the court questioned the citations, the attorneys compounded the failure by submitting AI-generated excerpts of the fabricated opinions. Judge P. Kevin Castel found subjective bad faith, imposed a \$5,000 sanction, and ordered the attorneys to notify each judge whose name had been falsely attributed to a nonexistent ruling. The court noted that nothing prohibits using AI for research, but the obligation to verify work product and maintain a reconstructible record remains with the professional. (Opinion and Order, *Mata v. Avianca, Inc.*, 22-cv-1461 (PKC), S.D.N.Y., June 22, 2023.)

Failure mode 3: Vendor AI capabilities deployed without due diligence or contractual coverage.

A product you already license adds AI processing. Your contract predates the feature, your diligence questionnaire does not ask about model providers or subprocessors, and the change arrives without a procurement event to trigger review.

IN PRACTICE · FAILURE MODE 3

A vendor you have used for years turns on an AI feature by default. Your regulated data is now processed by a subprocessor you have never assessed and cannot name, under a contract that does not mention it.

Cited incident. In August 2023, Zoom updated its Terms of Service to grant itself broad rights to use customer audio, video, and chat content for machine-learning and AI model training, with no opt-out mechanism at the time of the change. Customer data—including recordings of medical consultations and confidential business meetings—became subject to processing that existing contracts had never contemplated. Zoom later acknowledged the episode as a "process failure" and revised the terms on August 11, 2023, to require customer consent before using communications content for AI training. Separately, the FTC's consent order against Easy Healthcare Corporation (the developer of the Premom fertility-tracking app), finalized June 2023, imposed \$100,000 in civil penalties after the company shared sensitive user health data with third-party analytics firms—including two China-based companies—via embedded SDKs, under a privacy policy that promised it would not do so. The order required deletion of the improperly shared data and barred future sharing without affirmative consent. (Zoom Terms of Service revision, August 11, 2023; FTC, *In re Easy Healthcare Corp.*, Dkt. No. C-4791, consent order finalized June 2023, FTC press release, May 17, 2023.)

In each case, the failure is an absence of systemic control rather than a malicious actor. That distinction matters, because it means the remediation is structural and therefore tractable.

1.2 The Cost of Inaction vs. the Cost of Action

The enforcement landscape is moving in one direction. The EU AI Act introduces penalties reaching 35M EUR or 7% of global turnover for prohibited practices. U.S. state-level AI laws are proliferating on their own timelines. Meanwhile the regulations you are already subject to apply to AI contexts without needing to be rewritten: HIPAA governs PHI regardless of what software touches it, Section 5 of the FTC Act reaches deceptive AI claims, and state bar disciplinary rules reach the lawyer, not the tool.

There is also a cost to doing nothing, which organizations routinely leave out of the analysis. A blanket AI ban does not eliminate AI use; it encourages relocation to personal devices where you have no visibility at all. It also carries its own compliance exposure. Professional competence obligations across regulated fields do not come with a technology opt-out, and standards of care evolve toward what competent practitioners actually do.

Governed adoption costs real money and real attention. The three postures compare like this.

TABLE 1. THREE AI ADOPTION POSTURES COMPARED

	No adoption	Ungoverned adoption	Governed adoption
Up-front cost	None	None	Policy development, contract review, training, tool procurement
Ongoing cost	Manual process cost stays flat or rises	Invisible until it is not	Monitoring, periodic reassessment, governance committee time
Regulatory exposure	Competence and standard-of-care drift	Undocumented data flows, no audit trail, uncovered subprocessors	Documented, bounded, and defensible
Competitive position	Falling behind on price and turnaround	Uneven; some teams gain, others create liability	Capacity gains you can put in front of a client or a regulator
On a bad day	You explain why you had no capability	You reconstruct an incident with no records	You produce the log, the policy, and the contract

Each column represents an organizational posture toward AI adoption. The middle column—ungoverned adoption—is where most organizations sit today and where the risk-to-visibility ratio is highest.

The column that fails worst under examination is the middle one. Ungoverned adoption gets you the benefits and the liability at the same time, and it is where most organizations are sitting right now.

1.3 The ROI Framework

A working AI compliance infrastructure produces returns in three forms: work that becomes economically possible, capacity multiplication, and risk reduction.

Work that was previously cost-prohibitive. The speed figure is real. Document Review and Questionnaire response are substantially faster with AI assistance taking the first pass. But speed is the mechanism, not the return. The return is that the skill lays down a base-level analysis that can be confirmed faster than it could have been performed by hand, and produces a fully built-out artifact that can be reviewed, edited, and delivered to a client.

Producing that artifact by hand, for every change to every document, is something many orgs would choose not to do for every change, but with properly controlled AI tooling that is possible. Changes that used to be a judgment call and a note in an email are becoming tracked updates, and comprehensive analyses that were once annual activities are becoming constant monitoring activities.

The real return is not speed. It is that the cost of performing and documenting an analysis has dropped far enough that you can run it on every event, not just the ones large enough to justify the hours.

The same shift applies to any task where mechanical work dominates and the output is a deliverable: contract clause extraction, policy-to-framework mapping, regulatory change triage, vendor questionnaire review responses.

Capacity multiplication follows from the same change. When human time shifts from producing the analysis to confirming it, throughput changes shape. A consultant can carry several times the client load. A compliance department can monitor several times the vendor population. The practitioner's judgment is not removed from the process; it is moved to the point in the process where it is worth the most.

Risk reduction One of the sources of risk reduction comes from closing the gap between when something changes and when you find out. An annual review cycle leaves a twelve-month exposure window by design unless the change is detected during periodic reviews, smaller scoped risk assessments, or continuous monitoring. If a regulation changes in month two, you could carry the gap for ten months. Continuous monitoring narrows that window and can allow the organization to respond to the change before waiting for a periodic manual analysis for triage and action.

What "AI-Ready" Actually Means for a Regulated Organization

"AI-ready" gets used as though it were a single condition. It is set of conditions that are a combination of technical and non-technical requirements, and they must all be met for AI tooling to succeed.

Readiness can be broken down into four dimensions: data, documents, governance, and infrastructure. Each section below carries a scored instrument, and each one resolves to the same four bands: Absent, Reactive, Defined, Managed. Run all four, record them on the scorecard at the head of Section 2.5, and that section turns the results into a sequenced plan.

2.1 Data Readiness

Your data needs to be in a format AI can consume.

On the regulatory side, federal regulations are machine-accessible. The eCFR publishes an API, and a pipeline that pulls eCFR content into a structured knowledge base is straightforward engineering. State regulations are a patchwork: different formats, different update cadences, different access models, and in some cases no programmatic access at all. Each state is its own integration project. Most organizations have never structured their own regulatory applicability at all, which means the first question any AI system asks, "which rules apply here," has no stored answer.

On the organizational side, the pattern is familiar. Contracts sit in filing cabinets. Policies live in shared folders with no version control. Training records are in spreadsheets. The vendor list is in someone's email, and there is a second vendor list in procurement that does not match.

A compliance data pipeline means regulatory and organizational data flow into one system that can reason about both at once. The requirement is a system that knows which rules apply to you and what you already have in place against them.

TABLE 2. DATA READINESS INSTRUMENT

Data Readiness Self-Assessment

Score each question 1 to 4, then total and read the band.

#	Question	1	2	3	4	Score
1	Can you produce a complete list of regulations that apply to your organization?	No, it would take a project	Days, assembled by hand	Hours, from existing documentation	Under an hour, from a maintained source	
2	Are your contracts, policies, and procedures stored in formats AI can read?	Paper and scanned images	Mixed, mostly scanned PDFs	Mostly text-based files	Text-based and consistently structured	
3	Do you have a single source of truth for your vendor and Business Associate list?	No list exists	Multiple lists that disagree	One list, updated periodically	One list, updated on contract events	
4	When a regulation that applies to you changes, how do you find out?	An auditor, client, or regulator tells you	Someone notices during other work	A subscription or newsletter routed to a named owner	Monitoring that maps the change to your affected policies	
5	Could a new compliance hire understand your current posture without asking someone?	No	Only for part of the program	Mostly, with some tribal knowledge	Yes, the documentation is the source of truth	

Your data readiness total _____ / 20

Score each row 1 through 4 based on where your organization falls today. A score of 1 indicates the lowest maturity; 4 indicates the highest. Total your scores and match to the bands below.

SCORING BANDS (5 TO 20 POINTS)

Band	Score	What it means
Absent	5 to 9	Data work is the prerequisite. Start at Phase 1; AI tooling deployed now produces confident output over incomplete inputs.
Reactive	10 to 13	Enough structure to pilot narrowly. Scope any deployment to the one area where your data is cleanest.
Defined	14 to 17	Ready for scoped deployment across multiple areas. Focus remediation on your lowest-scoring question.
Managed	18 to 20	Ready for continuous monitoring. The constraint is now governance, not data.

2.2 Document Readiness

Most organizational documents were written for human readers, and human readers are forgiving. They tolerate a policy with three revision dates in the header and no indication which one is current. AI systems are not forgiving in that way. They will read all three and treat every one as authoritative. This leads to inconsistent results and unreliable analysis. These documents need to be current, version controlled, and stored in formats that AI can read.

Document readiness has two components: format and governance. Format determines whether a system can read the document at all. Governance determines whether reading it produces the right answer.

The compliance document taxonomy. These are the document types that feed a compliance AI system, in rough order of how often they are needed:

Document type	What the system does with it	Common blocker
Policies and procedures	Maps organizational commitments against regulatory requirements	No version control; multiple copies in circulation
Contracts, BAAs, and DPAs	Extracts clauses, flags missing required provisions, tracks subprocessors	Scanned signature copies; every vendor uses a different template
Training materials and completion records	Evidences workforce training obligations	Records live in a learning system that does not export
Incident logs	Establishes response history and pattern detection	Kept in email threads rather than a log
Vendor and Business Associate assessments	Feeds due diligence and continuous monitoring	Point-in-time, never refreshed
Risk assessments and audit reports	Establishes baseline posture and open findings	Produced as final-form PDFs with no underlying data
Board and committee minutes	Evidences governance oversight	Access-restricted in ways the system cannot accommodate

The document quality spectrum. Where a document sits determines what it costs to make usable:

Level	Description	Remediation
Unusable	Paper, handwriting, photographs of documents	Manual re-entry
Poor	Scanned image PDFs with no text layer	OCR plus human verification pass
Workable	Text-based PDFs and word processor files	Extraction, generally automatable
Good	Consistent text formats with reliable headings	Light structuring
AI-optimized	Structured markup, tagged metadata, version-controlled	None; this is the target

TABLE 3. DOCUMENT READINESS INSTRUMENT

Document Readiness Self-Assessment

Score each question 1 to 4, then total and read the band.

#	Question	1	2	3	4	Score
1	Do your policies extract as text rather than as an image?	Paper or scanned images only	Mixed, mostly scanned	Mostly text-based	Text-based and consistently structured	
2	Does every policy have one current version, in one location, with a named owner?	Multiple copies circulate	One location, ownership unclear	One current version and a named owner for most	For all, with superseded versions marked as superseded	
3	Do your contracts follow a consistent structure?	Every vendor's own template	A house template applied inconsistently	House template for all new contracts	House template, with legacy contracts mapped to it	
4	Do you know when each policy was last reviewed?	No	For some	For all, tracked manually	For all, with the next review date scheduled	
5	Can you export from your document, learning, and contract systems without a vendor ticket?	No	From one of the three	From most	From all, on demand	

Your document readiness total

/ 20

Score each row 1 through 4 based on where your organization falls today. A score of 1 indicates the lowest maturity; 4 indicates the highest. Total your scores and match to the bands below.

SCORING BANDS (5 TO 20 POINTS)

Band	Score	What it means
Absent	5 to 9	Start at Phase 1. Document remediation belongs ahead of any tool selection.
Reactive	10 to 13	Usable for a narrow pilot. Pick the document set that scored best and scope to it.
Defined	14 to 17	Documents will support analysis across multiple areas. Close your lowest-scoring question.
Managed	18 to 20	Documents are an asset rather than a blocker.

2.3 Governance Readiness

Many organizations are formalizing an AI policy while pilot projects are in progress or concluding. Tools arrived before governance did, and the governance is being written after the first deployment rather than before it.

An AI governance framework contains four artifacts. They are separable, and most organizations start by building these four artifacts.

ARTIFACT 1

AI Acceptable Use Policy

What employees can and cannot do with AI tools, which data classifications can be input to which tools, the approval path for adopting a new tool, and the consequence of going around it. The policy that fails is the one that only says no; staff route around it and you lose the visibility the policy was supposed to create.

ARTIFACT 2

AI Vendor Due Diligence Process

How you evaluate a vendor's AI capabilities, which contractual protections are required before regulated data touches the product, and how far down the subprocessor chain you require visibility. This process has to fire on existing vendors adding AI features, not only on new procurements. That is the trigger most processes miss.

ARTIFACT 3

AI Incident Response Protocol

What happens when an AI-influenced output reaches a client, patient, or regulator and turns out to be wrong. How you trace which system produced it and on what inputs. Who owns the correction and who owns the notification analysis. This should extend your existing incident response plan rather than sitting beside it as a separate document.

ARTIFACT 4

AI Governance Committee Charter

Who owns AI governance, what they decide without escalating, where escalation goes, and how often they meet. Ownership that lives nowhere in particular defaults to whoever was in the room when the question came up.

TABLE 4. GOVERNANCE READINESS MATRIX

Those four artifacts are the four rows below. Find your organization in each row: the stage you sit at is that artifact's maturity, and the row you score lowest on is usually where the next incident comes from.

Dimension	Stage 1: Absent	Stage 2: Reactive	Stage 3: Defined	Stage 4: Managed
Acceptable Use	No policy	Informal guidance issued after an incident	Written policy, communicated to the workforce	Policy enforced, monitored, and updated on a cycle
Vendor Diligence	Not considered	Ad hoc review when a question arises	Standardized questionnaire applied at procurement	Continuous monitoring, including existing vendors adding AI
Incident Response	No plan	Handled case by case	Documented protocol integrated with existing IR	Tested, drilled, and improved from findings
Committee and Ownership	Nobody owns it	Assigned to someone after a problem	Standing committee with a charter	Integrated into existing governance structures

Your governance band is your lowest stage across the four rows, not the average. Governance fails at its weakest artifact.

Your governance band _____

A governance band of Absent starts you at Phase 1 regardless of the other three dimensions. Your governance band is your lowest stage across the four rows, not your average. Governance fails at its weakest artifact, and averaging hides that.

Stage 3 across all four rows is the realistic target for a first year. Stage 4 in one row and Stage 1 in another is a common and unstable pattern; it usually means one motivated person built the artifact they cared about and the rest did not follow.

2.4 Infrastructure Readiness

Three questions carry this entire section: where does the AI run, who controls it, and what data does it touch. All three are business questions with contractual answers, which is why they belong to whoever owns the regulatory obligation.

The three deployment models

	Public cloud AI	Private or on-premise AI	Hybrid
Examples	Anthropic, OpenAI, and Google APIs; vendor products built on them	Local models run through Ollama or vLLM on hardware you control	Local models for routine work, cloud APIs for high-stakes analysis
Where data goes	Leaves your environment	Stays inside your environment	Routed by data classification and task
Contract required	depends on deployment	None; you are the processor	Required for the cloud path only
Strengths	Best available model quality; no infrastructure to maintain	No third-party data exposure; no vendor dependency	Sensitive data stays local without giving up model quality where it matters
Trade-offs	Vendor dependency; contractual and subprocessor diligence burden	Requires technical capability; model quality gap on hard reasoning tasks	Two environments to govern; routing logic has to be explicit and auditable
Fits	Organizations with mature vendor management and non-sensitive or contractually covered use cases	High-sensitivity data, strong internal IT, tolerance for narrower capability	Organizations with both the vendor management maturity to handle frontier model deployments and the technical capability to build and maintain local solutions

The "where does the data go" decision tree

Run this per use case, not per organization. Most organizations need more than one answer.

1 Does the task touch regulated data?

No Any model is acceptable subject to your acceptable use policy. **Stop here.**

Yes Continue to 2.

2 Can the data be de-identified or redacted without destroying the task?

Yes De-identify the data, then return to question 1.

No Continue to 3.

3 If you are subject to HIPAA, do you have a signed BAA or DPA with the model provider (or covering subprocessors)?

No Execute one OR route this use case to a local model.

Yes Continue to 4.

4 Would a wrong output reach a client, patient, or regulator without a human reviewing it first?

Yes This use case requires a documented human review step before deployment, regardless of model.

No Continue to 5.

5 Can you reconstruct, six months from now, what the system was given and what it returned?

No You have a logging gap, not a model gap. Close it before deployment.

Yes **Proceed with deployment.**

Question 5 is the one that gets skipped, and it is the one that determines whether you can defend the deployment later.

Integration points. AI tools deliver very little when they sit beside your systems rather than inside them. These are the connection points that matter, and the governance question each one raises:

System	What the AI does there	Governance question
Document management	Reads policies and procedures for mapping and gap analysis	Can the system distinguish versions and respect the progression from draft to current to superseded?
Contract management	Extracts clauses, flags missing provisions, tracks renewal and subprocessor changes	Who reviews AI-flagged findings before they become positions?
Compliance management	Maintains framework mappings and open findings	Is AI-generated evidence labeled as such in the audit record?
Email and calendar	Surfaces obligations and deadlines from correspondence	Does this expand the scope of what leaves your environment?
Ticketing and incident systems	Triages and routes compliance issues	How is automated triage verified—periodic review, sampling, or full human review—and who owns the catch rate?

TABLE 5. INFRASTRUCTURE READINESS INSTRUMENT

Infrastructure Readiness Self-Assessment

Score each question 1 to 4, then read the band. The bands match Section 2.1 so the four dimensions are comparable.

#	Question	1	2	3	4	Score
1	For each AI use case in production, do you know where the data goes?	No use cases are documented	Known for some	Known for all current use cases	Known and recorded as part of the approval path	
2	Is every cloud path that touches regulated data covered by a signed BAA or DPA?	No agreements in place	Some, coverage unverified	All direct providers covered	All providers covered, with subprocessor terms verified	
3	Is routing between local and cloud models explicit?	No local option exists	Left to individual choice	Guidance exists but is not enforced	Enforced by configuration, not by instruction	
4	Does every output that can reach a client, patient, or regulator pass a documented human review step?	No review step defined	Informal review by whoever is available	Documented reviewer per use case	Documented, with a defined path for reviewer disagreement	
5	Could you reconstruct, six months from now, what a system was given and what it returned?	No logging	Partial logs, no retention policy	Inputs and outputs logged and retained	Logged, retained, and retrievable by case within your evidence process	

Your infrastructure readiness total _____ / 20

Score each row 1 through 4 based on where your organization falls today. A score of 1 indicates the lowest maturity; 4 indicates the highest. Total your scores and match to the bands below.

SCORING BANDS (5 TO 20 POINTS)

Band	Score	What it means
Absent	5 to 9	Any current AI use is undocumented and undefendable. Start at Phase 1.
Reactive	10 to 13	Deployment is possible for the use cases you can fully account for, and only those.
Defined	14 to 17	Ready to expand deployment. Close the lowest-scoring question before you do.
Managed	18 to 20	Ready for continuous operation. Logging and routing hold up under examination.

2.5 From Assessment to Plan: The Four Phases

Four assessments produce four bands. The bands are only useful if they sequence into work. This is the phase model the vertical supplements build on, and the ordering is deliberate: each phase produces the inputs the next one consumes.

TABLE 6. READINESS SCORECARD

Readiness Scorecard			
Dimension	Section	Score	Band
Data	2.1	_____ / 20	_____
Documents	2.2	_____ / 20	_____
Governance	2.3	Stage _____ / 4 (enter your lowest stage from Table 4)	_____
Infrastructure	2.4	_____ / 20	_____

Transfer your scores from each instrument above. Your lowest band sets your starting phase.

Your lowest band sets your starting phase. An organization at Absent on any dimension starts at Phase 1 regardless of how the other three scored, because the phases depend on inputs the weak dimension has to supply.

PHASE 1 **Inventory** Find out what is actually happening.

- Produce an AI tool inventory covering what the workforce uses today, sanctioned or not. Ask without penalty attached, or the inventory will be wrong.
- Map data flows: what data classes exist, where they live, and which of them have already touched an AI system.
- Produce the regulatory applicability list from Section 2.1, question 1.
- Inventory vendors and Business Associates, including which of them have added AI features to existing products.
- Score all four readiness dimensions and record the baseline. All four resolve to the same four bands, so record the band and the underlying score together.

Exit criterion: you can state what runs where, on what data, under which contract. Where you cannot, you have written down that you cannot.

PHASE 2 Foundation Build the governance scaffolding before expanding use.

- Draft and adopt the four governance artifacts from Section 2.3.
- Amend contracts where the diligence process found gaps, prioritizing vendors that already process regulated data.
- Remediate the document gaps that scored lowest in Section 2.2, starting with the document sets Phase 1 found already in use.
- Train the workforce on the acceptable use policy, including the approval path. A policy nobody can follow procedurally is a policy nobody follows.
- Stand up the governance committee and hold the first meeting.

Exit criterion: an employee with a new AI use case has a documented path to approval, and the person who approves it exists.

PHASE 3 Deployment Put approved tools into production against real work.

- Select use cases using the Section 2.4 decision tree, starting where data sensitivity is lowest and the manual cost is highest.
- Run a pilot with success criteria defined before it starts, including an accuracy baseline measured against current manual output.
- Implement the safeguards the decision tree required: review steps, logging, routing.
- Document the human review step for each deployed use case. Who reviews, against what, and what they do when they disagree.

Exit criterion: at least one use case runs in production with logging, a named reviewer, and measured results against the pre-pilot baseline.

PHASE 4 Operations Move from project to standing capability.

- Turn on continuous regulatory monitoring and route changes to the policies they affect.
- Set the reassessment cadence for risk assessments and vendor reviews, and tie it to change events rather than only to the calendar.
- Drill the incident response protocol at least once before you need it.
- Package audit evidence continuously rather than assembling it under deadline.
- Re-score the four readiness dimensions annually and compare to the Phase 1 baseline.

Exit criterion: the compliance posture is a maintained state rather than a periodic reconstruction.

Phases 1 and 2 are where organizations most often try to save time, and they are the two that determine whether Phases 3 and 4 work at all.

2.6 Building Your Readiness Plan

The scorecard tells you where you are. The phases tell you what order to work in. This section helps you turn both into a document you can bring to leadership, scope a budget against, or hand to an outside partner.

TABLE 7. READINESS PLAN TEMPLATE

Readiness Plan Template	
Element	Your entry
Starting phase (from Table 6, lowest band)	_____
Weakest dimension (lowest-scoring assessment)	_____
Lowest-scoring question in that dimension	Q _____ : _____
Second-weakest dimension	_____
Lowest-scoring question in that dimension	Q _____ : _____
First remediation target (the gap that blocks the most downstream work)	_____
Owner (name and role, not a department)	_____
Target date for Phase 1 exit criterion	_____
How you will know it worked (measurable outcome, not effort)	_____

Complete this template using your scores from Tables 2 through 6. The lowest-scoring question in your weakest dimension is usually the first thing to fix, because it gates the work that follows.

The plan does not need to be long. One page that states where you are, what you are fixing first, who owns it, and when you will know it worked is more useful than a thirty-page strategy deck that nobody executes. If you cannot fill in the Owner row with a person's name, that is your first problem to solve.

The Orchestration Layer

The gap is not capability. The models are genuinely good at reading, reasoning, and following instructions. The gap is that nobody gave them the right instructions.

This section is the honest accounting of what stands between a capable model and a working compliance system.

3.1 Why Data Sources Are Not Ready Today

Federal regulations are machine-accessible but not machine-ready. The eCFR API returns the text. It does not return which sections apply to your organization, how they relate to the framework you report against, or which of your policies address them. That mapping is engineering plus judgment, and it has to happen before a model can do anything useful with the text.

State regulations are worse in a specific way. Access, formatting, and update cadence vary by state, so there is no single approach that covers all of them. Every state you operate in is a separate integration with separate maintenance.

Organizational data has the digital attic problem. Most organizations hold years of policies, contracts, and records spread across structured databases, shared drives, a document management system that was migrated twice, and email attachments. The material is all there. It is not in one place, it is not consistently current, and nothing indicates which copy governs.

Remediating this is not a software install. It is a structured project with human judgment at every stage: deciding which copy is authoritative, which policy supersedes which, which regulatory sections actually apply. No model makes those calls for you.

3.2 The Compliance Integrator Role

Somebody has to sit between the regulatory requirement and the system configuration. That work means translating requirements into structures a system can act on, mapping organizational data to compliance frameworks, and designing workflows that encode practitioner judgment rather than prompting for it fresh every time.

At most organizations this work is already happening in pieces, without a name and without time budgeted for it. The compliance officer who built the vendor tracking spreadsheet that the whole department actually uses is doing it. So is the analyst who worked out which fields the reporting system needed before anyone had specified them. What is usually missing is the mandate, the time, and a counterpart on the technical side who understands what a regulatorily correct answer has to look like.

That is where the work can stall. Operations, Legal, or Compliance may hold the judgment about what a correct answer is depending on the subject matter. Judgment is the missing layer. AI models fail at compliance work not because they cannot read the regulation but because they lack the accumulated judgment about what matters in it. That judgment is not a prompt. It is the thing a practitioner builds over years and has never had to write down.

There are three ways to fill the role.

Build internal capability

Hire the profile or upskill someone who already has half of it, usually a compliance analyst with technical aptitude. This produces the most durable capability and the slowest timeline. Budget for a year before the capability is self-sustaining, and treat their subsequent advancement as a required part of the plan.

Engage outside expertise

Bring in a compliance integrator on a project basis to produce the governance artifacts and the initial data structuring. Faster, and it front-loads the judgment where it matters most. The failure mode is the deliverable that arrives as a document nobody operates; guard against it by requiring that the engagement leave behind a running process, not a PDF.

Adopt a platform that encodes the methodology

Buy the judgment already written down. Fastest to a working state and the least flexible; you are adopting someone else's model of how the work is done, which is fine when their model matches your work and expensive when it does not.

These are not mutually exclusive, and the common sequence is engage, then build, with a platform running underneath. Section 3.3 gives the signals that indicate which one to start with.

3.3 The Standard Your Peers Are Moving Toward

When continuous compliance monitoring works, four things are true. Regulatory changes surface within days rather than at the next annual review. Policy gaps are identified against the changed requirement automatically. Vendor compliance status is tracked live rather than reconstructed at renewal. Audit evidence is packaged as it accumulates rather than assembled under deadline.

That is not a distant target. It is the level leading compliance departments and consultants are building toward now, and the organizations that get there first will set the expectations that clients and regulators apply to everyone else.

For consultants specifically, the pressure arrives from a different direction. Your clients will start asking why you do not have this. Not because they read this guide, but because a competitor will offer it and they will notice the difference in turnaround.

Build-vs-Engage Decision Framework

Read down the column that most of your signals point to. Mixed results are normal and usually mean a sequenced approach.

Signal	Build internally	Engage consulting	Adopt a platform
Regulatory complexity	Single framework, stable	Multi-framework, evolving	Multi-framework, continuous change
Internal technical capability	Strong IT and compliance alignment	Compliance expertise, limited technical	Limited on both
Timeline pressure	12+ months acceptable	Governance needed in 90 days	Monitoring needed now
Budget profile	Can invest in headcount	Project-based budget	Subscription budget
Readiness scores from Part 2	Defined or Managed across dimensions	Reactive, with clear gaps to close	Absent, needing structure before anything else

3.4 What Comes Next

This guide covered the organizational and strategic framework. The Healthcare Supplement and the Legal Supplement apply it to specific regulatory surfaces and implementation playbooks.

If you want to know where you actually stand, run the four assessments in Part 2 and record the results. That costs you an afternoon and no money, and it tells you which phase in Section 2.5 you are actually starting from.

NEXT STEP

Dang's Solutions offers AI governance assessments that produce the artifacts described in Section 2.3. If this sounds like a fit, I would like to talk. We will go through what you are working on and I will tell you honestly whether it is something I can help with.